

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

# 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

Idea generation is easy. We could brainstorm new product ideas all day. The hard work comes when validating a new idea improves the product. A/B testing can help validate new ideas so that we're improving the product for both the users and the business. To do this, we need data to demonstrate the impact of an A/B test.

You may have noticed an important detail yet to be discussed, which is the complexities of gathering, querying, and visualizing data for your experiments while also ensuring your test configuration is valid.

The previous chapter focused on the different types of experiments. We dabbled in non-inferiority tests that seek to measure if a change is as good as the control, given a predefined margin. We learned the logistics for evaluating changes on a longer time scale by leveraging long-term holdbacks for a degradation or cumulative holdback test.

Regardless of which type of A/B test is selected, you'll need access to quality data to analyze and adequately visualize your test results. How you interface with data to create your A/B test results directly influences the A/B testing process. In this chapter, we'll explore the following:

- Common practices for storing and accessing data.
- Different types of data to create product insights and monitor the performance of an A/B test.
- Various visualizations to present the results of an A/B test.

Let's tackle the more complex and data-intensive aspects of A/B testing. Here we go!

## Working with Different Types of Data

The goal of A/B testing is to observe what happens as you introduce changes to the product. To do this, you need data, a lot of data. You also need different types of data. Having different categories of data will help you with the following:

- Enable you to cross-reference and validate outcomes to increase confidence in the analysis.
- Give you the capability to monitor your engineering system health that creates the product experience.
- Provide you with a deeper, richer context for data analysis.

Let's look at the different categories of data that will create a deeper understanding of how your changes impact the product.

## Defining User Engagement Data

To understand what users are up to on the product, you'll need data representing their actions and engagement. The typical path to get user engagement data is for the front end, or client side, to log events. The events will initially be raw and maybe a bit messy. When aggregated and enriched, these events can be used to understand the user's journey as they engage with the product. Examples of user engagement events include the following:

- Impressions
- Clicks
- Pageviews

When you have user engagement events such as impressions, the work to create the events is referred to as instrumentation. Let's emphasize the importance of instrumentation. To run an experiment, you'll need instrumentation within the product to create the raw events, or logs, representing how users engage with product features. Instrumentation is a crucial mechanism that will enable you to understand what the users see on the product and then do as a result of seeing a change or new feature. If you're unsure what your product should instrument, think about the actions a user can take on the product. Can a user click or scroll? How long does it take them to scroll or click something they recently saw? Implementing instrumentation will give you a richer understanding of the effect of changes introduced in an experiment.

Remember that these data points alone are usually meaningless as a unit by themselves. Any data, regardless of the type, is only valuable if you can understand it. Therefore, enriching and aggregating it to improve interpretation is key. Let's look at the data architecture at CableMax to illustrate how user engagement data can be combined with additional contexts to create a richer understanding.