

Vignette® Measuring Recommendation Impact

Product Datasheet

Measuring the impact of Recommendations and Social Search on a Web site is often critical to proving project success. As such, Vignette offers a robust testing platform that can meet any number of testing scenarios. The two most common scenarios are described below:

Testing scenario 1: competitive A/B split test on directed revenue lift

It is sometimes the case that a customer already has an existing recommendation strategy deployed on the site, such as manually merchandized recommendations or a basic “also bought” algorithm. If the customer wants to test the effectiveness of Vignette Recommendations versus their existing recommendations, an A/B split test is most appropriate. In such a test, half of all users will see Vignette recommendations (the A group) and the other half see the pre-existing recommendations (the B group). Once a user is assigned to a group, that user will continue to be a part of that group for the duration of the test—even across multiple sessions and multiple days or weeks.

There are several ways in which to measure and compare the effectiveness of the two test groups. The most direct and noise-resistant method is based on **directed revenue**. Directed revenue considers only those users who bought exactly what was recommended to them. This is the clearest way to ensure causality between the recommendations offered and subsequent purchases.

For example, let’s say a user sees a recommendation that includes three products. If that user clicks on the first product and then purchases it—either immediately or within a specified time window—the recommendation is given credit for that purchase and the revenue associated with that product is considered to be directed revenue. If that user added other things to the cart beyond this product, that additional revenue would not be considered directed revenue unless those products were also preceded by a recommendation click on those specific products.

Products that are purchased without direction from a recommendation are not included in the comparison. So users who purchase in either condition without seeing or using recommendations, or buy things other than what was recommended are not considered in the comparison. Non-directed purchases can be very noisy and prevent having a clear picture of the benefit on recommendation strategy offered one over another.

Below is an example table showing what results from an A/B Split Test based on directed revenue might look like:

	A	B	Lift
Total Directed Revenue	\$30,000	\$10,000	200%
# Purchases	150	60	150%
Recommendation Users	2,000	1,000	100%
Revenue Per User	\$15	\$10	50%
Conversion Rate	7.5%	6%	25%
AOV	\$200	\$167	20%

Testing scenario 2: ongoing test on influenced revenue lift

The second most common test scenario is used when a customer does not have an existing recommendation strategy to compare to, but wants to measure the impact of Vignette recommendations on the site versus nothing. In fact, it is important to maintain such a measurement strategy for as long as recommendations are deployed so that the ongoing value can always be seen. In an ongoing test, all users are shown recommendations and thus have the opportunity to use them. The two groups for comparison then become those who use recommendations and those who do not.

Note that in such a test, it is likely that there will be more non-users of recommendations than recommendation users, depending on the exposure of recommendations and other factors such as purchases from telesales reps, repeat orders from non-shopper buyers. It is, however, meaningful to look at what percentage of total revenue is influenced by recommendations, as well as to directly compare metrics such as conversion rate, AOV, and revenue per user between the two groups.

Unlike the split test scenario, where measuring directed revenue is most appropriate, there is no notion of directed revenue in the non-users case to compare to. Instead, the fair comparison is on **influenced revenue**. In other words, recommendations will get credit for the purchase as long as any recommendation was used within a specified time prior to the purchase.

Below is an example table showing what results from an ongoing test might look like:

	Recommended Users	Non-Users	% of Total
Total Revenue	\$40,000	\$60,000	40%
# Purchases	300	700	30%
# Users	2,500	7,500	25%
Lift			
Revenue Per User	\$16	\$8	100%
Conversion Rate	12%	9%	33%
AOV	\$133	\$85	57%

Find Out More About Vignette Measuring Recommendation Impact

Vignette Measuring Recommendation Impact is part of Vignette's Web Experience Platform, an integrated family of enterprise content management products. With these products, Vignette's customers improve their content-driven business processes and deliver any content, to both internal and external audiences, anytime, anywhere and to any device making the Web experience relevant, social and multi-channel. For more information about Vignette's products, solutions and their demonstrated value to organizations in virtually every industry, please visit www.vignette.com or call +1 888 608 9900.