



Case Study: Landing Page Test

Big test, fast results, large ROI

Introduction

The Internet is an ideal marketing channel for testing. Catalog and retail tests have some limitations due to higher production costs and data lags. But with the speed, flexibility, and low cost of e-mail, online advertising, and landing page tests, the Internet is an excellent place to try new things and prove what works. Plus what proves effective online can then be tested and implemented in off-line programs.

One market leader saw impressive results after testing numerous changes to their landing page—the separate page people see, before the “official” website, when they click on a link from an e-mail, partner site, or online ad. Scientific testing let them:

- **Test 23 marketing-mix elements simultaneously over 3 weeks**
 - Split-run tests would have required 15 months for equal statistical confidence
- **Pinpoint and quantify 11 significant elements**
 - Six changes that increased conversion rate
 - Five “good” new ideas that hurt
- **Increase conversion rate 42% after implementing the optimal combination**

#1: Creative Freedom

The vice president of a specialty-clothing firm challenged his team to use more data and less opinion. “Guesswork is like sailing against the wind,” he said. “You tack back-and-forth hoping to make headway slowly and inefficiently. Testing is like sailing with the wind at your back—running free straight ahead as fast as your ship can sail.”

Split-run testing had delivered some good results, but he remained impatient. Improvement came slowly testing just one variable at a time. After researching alternatives, he turned to scientific testing and found an experienced “captain” to guide his team.

In the first meeting, the consultant introduced the new approach and led a brainstorming session. Starting with their current landing pages and looking at past tests and competitive sites, the team brainstormed 99 changes that might increase conversion rate and profitability.

#2: Scientific Discipline

The next step was to concentrate the power of their ideas. With guidance, the team narrowed the list to 23 elements, focusing on bold, independent, and actionable changes.

The team then chose two levels to test for each. The “minus” level was usually the control. The “plus” level was a change that someone thought would increase conversion rate or profitability. Multiple levels can be tested for each element, but two *bold* levels are usually most efficient when testing a large number of new ideas. The few most important elements can then be tested in greater detail.

Details of the 23 test elements (below) will be discussed along with the results of the test.

Test Element	(-) Control	(+) New idea
A Headline focus	Image (general message)	Product (specific item)
B Main product	End of season product	New product
C Main product picture	Product alone (plain)	Lifestyle (on location)
D Size of main picture	Control	Larger
E Number of secondary products	3	6
F Location of secondary products	Below	Side
G Price points	Mixed	All low
H Sale price format	Show "regular price"	Sale price alone
I Copy	Control	Less
J Special savings starburst	Small	Large and brighter
K Shopping bag graphic on header	No	Yes
L Links	Text link + small button	Text link only
M Order button ("add to shopping bag")	Control	Larger
N Cross-sell message	No	Yes, push accessories
O Credit card message	No	Yes
P Guarantee message	Yes	No
Q Navigation bar on side	Numerous links	Few links
R Phone number	Control	More prominent
S Sale items link	None	Separate link on navigation bar
T New products link	Control	Highlight
U Page two	Many products	Few products
V Order form	Control	Numbered steps
W Special-offer pop-up	No	Yes (save 10% with e-mail sign-up)

Since only landing pages were tested, the company’s home page was not affected. The company’s server was programmed to randomly show one test page to every visitor who clicked through partner sites, online ads, and e-mails. With 70-75,000 landing page impressions expected per week, the test could be completed in just two weeks, but the team planned to extend the test into a third week in case week-to-week differences affected results.

With an average conversion rate of 8.4% for the control landing page and about 72,000 impressions per week, the team “captain” calculated a good chance (80% likelihood) of seeing any element that changed conversion rate by at least 4.9% and a fifty-fifty chance of seeing effects as small as 3.4%. Split-run tests would require the same 2-3 weeks for one variable alone!

Step #3: Statistical Power

The captain created a 24-recipe scientific test design (below). Each of the 24 versions of the landing page included a unique combination of plus and minus levels of all 23 elements. Whereas each split-run test cell gives one data point on one variable, each scientific test recipe provides a new piece of information about *every* variable in the test. After receiving “recipe sheets” from the captain, all 24 recipes were created in just three days.

	Headline focus	Main product	Main product picture	Size of main picture	Number of secondary products	Location of secondary products	Price points	Sale price presentation	Copy	Special savings starburst	Shopping bag graphic on header	Links	Order button	Cross-sell message	Credit card message	Guarantee message	Navigation bar on side	Phone number	Sale items link	New products link	Page two	Order form	Special-offer pop-up	
Recipe	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	Conversion
1	+	-	-	-	-	+	-	+	-	-	+	+	-	-	+	+	-	+	-	+	+	+	+	7.78%
2	-	-	-	-	+	-	+	-	-	+	+	-	-	+	-	+	-	-	+	+	+	+	+	8.63%
3	+	+	-	-	-	-	+	-	+	-	-	+	+	-	-	+	+	-	+	-	+	+	+	7.82%
4	-	-	-	+	-	+	-	-	+	+	-	-	+	+	-	+	-	+	+	+	+	+	+	9.67%
5	+	+	+	-	-	-	-	+	-	+	-	-	+	+	-	-	+	+	-	+	-	+	+	9.57%
6	-	-	+	-	+	-	-	+	+	-	-	+	+	-	+	-	+	+	+	+	+	+	-	10.32%
7	+	+	+	+	-	-	-	-	+	-	+	-	-	+	+	-	-	+	+	-	+	-	+	9.43%
8	-	+	-	+	-	-	+	+	-	-	+	+	-	+	-	+	+	+	+	+	+	-	-	7.63%
9	+	+	+	+	+	-	-	-	-	+	-	+	-	-	+	+	-	-	+	+	-	+	-	10.53%
10	+	-	+	-	-	+	+	-	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	9.95%
11	-	+	+	+	+	+	-	-	-	-	+	-	+	-	-	+	+	-	-	+	+	-	+	9.63%
12	-	+	-	-	+	+	-	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	8.67%
13	+	-	+	+	+	+	+	-	-	-	-	+	-	+	-	-	+	+	-	-	+	+	-	9.60%
14	+	-	-	+	+	-	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	9.53%
15	-	+	-	+	+	+	+	+	+	-	-	-	+	-	+	-	-	+	+	-	-	+	+	8.85%
16	-	-	+	+	-	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	+	6.93%
17	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	+	-	-	+	+	-	-	+	9.35%
18	-	+	+	-	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	+	-	9.23%
19	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	+	-	-	+	+	-	-	6.90%
20	+	+	-	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	+	-	-	9.10%
21	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	+	-	-	+	+	-	9.02%
22	+	-	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	+	-	-	+	8.53%
23	-	-	+	+	-	+	-	+	+	+	+	+	-	-	-	-	+	-	+	-	-	+	+	9.00%
24	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	8.40%

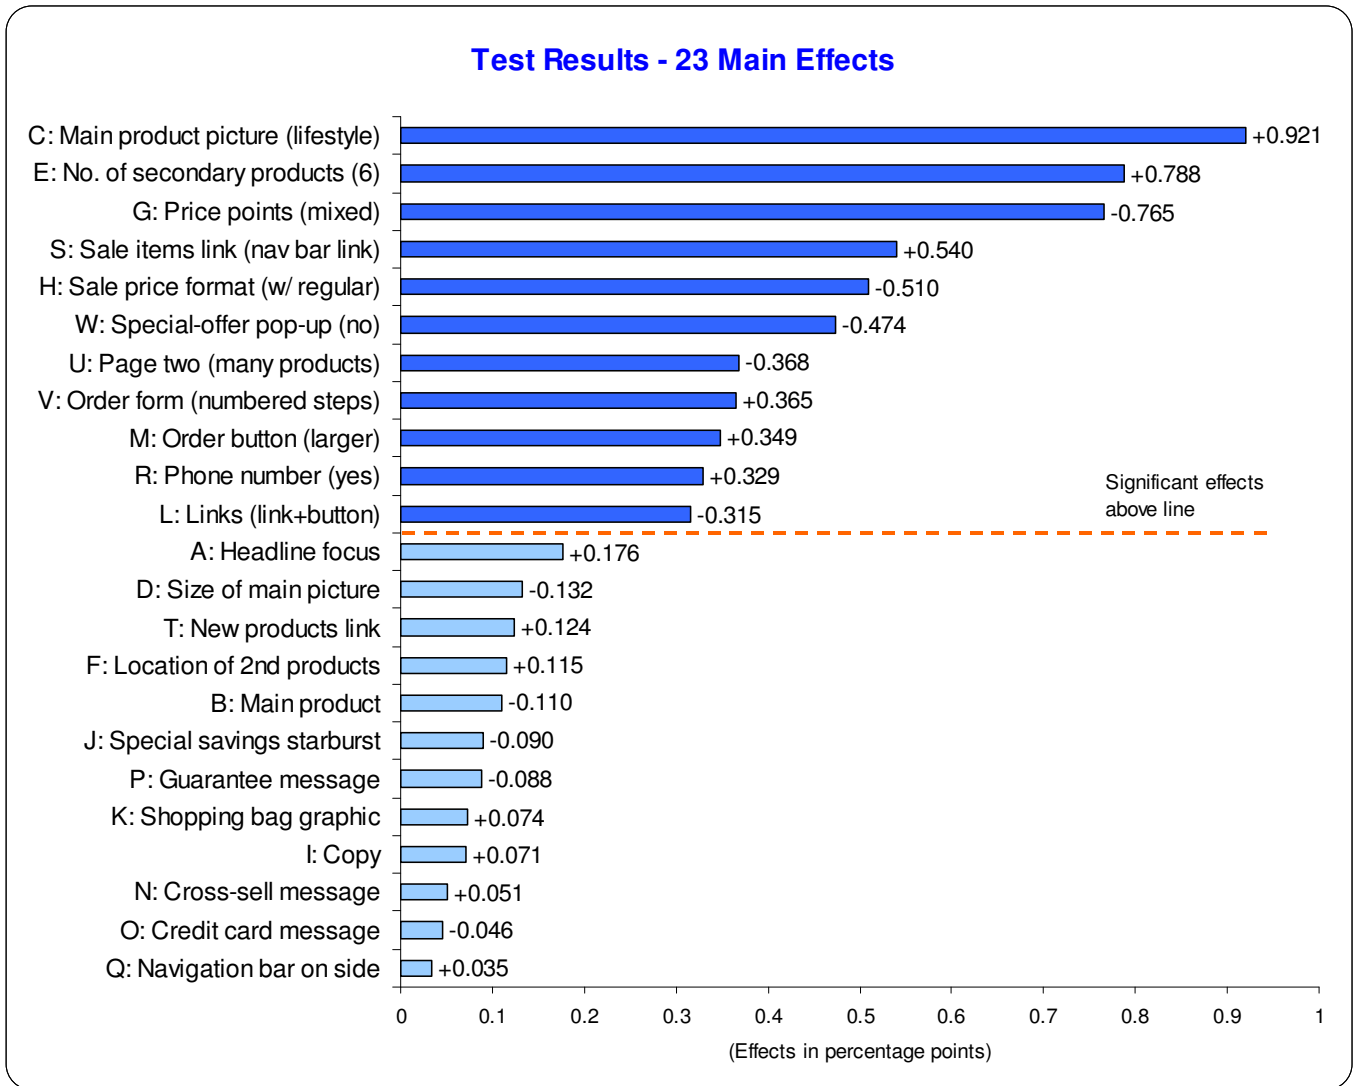
The key metric for the test was conversion rate—the percentage of people who view a landing page and then place an order. Average order size, gross margin, and differences among placements were also analyzed.

Each visitor who clicked through a landing page “recipe” was tagged with the recipe number, so wherever they went in the website, any purchase was tracked back to the original landing page.

#4: Marketing Insights

With about 14,400 total impressions over two weeks and 12,845 orders, the overall average conversion rate during the two-week test was 8.92%—higher than the 8.4% control. (Week 3 data is not included, but analyses over all three weeks showed the same results.)

The optimal combination of all 23 elements (shown below)—including 11 significant effects—predicted a 40.33% increase in conversion rate, from the control conversion rate of 8.40% to 11.78%.



Main effects of all 23 test element are shown above. With this simple chart, the marketing team could easily see the size of each main effect, the optimal level of each, and which are statistically significant.

In the chart, effects are ordered from largest (at top) to smallest (at bottom). The element name is on the left (with the optimal level in parentheses) and the effect is shown on the right as the length of each bar and the number at the end. The sign of each effect shows which level is better: with a positive effect, the “+” level will increase conversion and with a negative effect, the “-” level will increase conversion.

Significant Effects

Out of the 11 significant effects, 6 changes increased conversion (the “new idea” was better) and 5 changes hurt (the “control” was better). The significant effects included:

C+: Main product picture (lifestyle)

The “lifestyle” picture background—showing a person wearing the product in a scenic locale—was much better than the plain image of the product with a white background. This one change increased conversion 0.921 percentage points (pp), from the control conversion rate of 8.4 to 9.321 percent (an 11.0% increase).

E+: Number of secondary products (6)

Along with the main product on the landing page, having six additional products (smaller images grouped together) was much better than three secondary products, increasing conversion 0.788 percentage points (9.4%).

G-: Price points (mixed)

The team thought that a larger selection of low-priced items might increase response, but they were wrong. A mix of price points had an average conversion rate 0.765 pp (9.1%) higher. After seeing these results, they theorized that a broader product mix offered more choices with a few options still available for price-sensitive shoppers.

S+: Sale items link (separate link on navigation bar)

Without increasing the number of sale items, the team found that having a more prominent “on sale” link increased overall conversion by 0.540 pp (6.4%). They reasoned that the link did not affect the majority of buyers, but gave price-sensitive customers a short-cut to the best deals.

H-: Sale price format (show “regular price” along with the sale price)

The majority of items were not discounted, but the few discounted items on the landing page and all sale items normally showed the “regular price” along with the sale price in red. Some people on the team thought that showing both prices was unnecessary, but the test showed that the control was 0.510 pp (6.1%) better.

W-: No special offer pop-up

A pop-up window offering a 10% discount on the first order after signing up for the e-mail newsletter proved harmful. The pop-up reduced conversion by 0.474 pp (5.6%).

U-: Page two (many products)

“Page two” meant most of the pages visitors can click into from the landing page. In addition to testing more products on the landing page itself, the team tested many versus few products on each of these pages. Conversion rate was higher with more products on each page (the “-” level was 4.4% better).

V+: Order form (numbered steps)

Trying to simplify the order process and reduce the number of people who abandon their “shopping cart” before check-out, the team created a simplified form with numbered steps. The new form was 0.365 pp (4.3%) better.

M+: Larger order button

A larger order button with a 3-D appearance increased conversion 0.349 percentage points (4.15%).

R+: Phone number (more prominent)

A toll-free number had been at the bottom of the navigation bar in a small font. With a different toll-free number, moved to the top of the navigation bar in a larger font, the team saw a small, but insignificant rise in conversion rate. However, when they added phone orders to online orders, total conversion rate was 0.329 pp (3.9%) higher, allaying the team’s concern that the prominent number might just shift orders to a higher-cost channel.

L-: Links (text link + small button)

All main text links on the landing page also had a small arrow graphic in front of the words. One person thought that the graphic was an unnecessary annoyance since the text link is underlined. But the test showed that the graphic is helpful. Conversion rate dropped 3.75% when the graphic was eliminated.

Interactions

After analyzing all 23 main effects, the “captain” analyzed potential interactions. Three statistical principles guided his search for interactions among the hundreds of possibilities:

1. Main effects tend to be larger than two-way interactions and higher-order interactions are very unlikely (effect dissipation),
2. Few, if any, interactions are usually significant (effect sparsity), and
3. Interactions tend to result from large main effects (effect heredity).

In addition, a good understanding of the test elements helps guide the search for interactions. Interactions often result from conceptually or physically related elements. For example, price and offer variables often interact, plus related elements like the main product picture, background, and size (elements B-D) may create interactions.

In this case, potential interactions (including C, E, and G) were carefully analyzed, but none were found to be statistically significant.

Non-significant elements

The 12 effects that were not significant also provided valuable information. The additional cross-sell and credit card messages (N and O) were kept off to keep the page clear and uncluttered (though the guarantee message (P) was kept on at the bottom). Since the new special savings starburst, new products link, and navigation bar had no impact, these three were kept at the control level. Other elements were also kept at the control and the team planned to run additional headline and copy tests later.

[#5: Profit](#)

After creating the optimal landing page and running it against the original control, conversion rate jumped to 11.95%, a 42% increase over the control.

The marketing vice president was amazed that in less than two months, they uncovered 11 important elements and sailed to an impressive 42% increase in conversion rate. “Like a strong wind blowing away the morning fog,” he said, “we could immediately see the marketplace more clearly, know exactly where to go, and sail full-speed ahead.”

And this one test was only the beginning of their journey...