

What Impact Do “Bad Respondents” Have on Business Decisions?

We quantified the risks—and you’ll find the results both surprising and alarming.

By Michael Conklin
Chief Methodologist

Concern About Data Quality

The quality of online survey data is a major concern—not only for the market research industry but also for business leaders who rely on such data to make decisions.

Critical questions about data quality, however, have been difficult to answer. Can a small percentage of “bad respondents”—those whose responses are unreliable—affect survey results? Can the risk imposed by bad respondents be mitigated by increasing the sample size? How can data quality be assured when multiple panel sources are involved?

Recent studies by MarketTools quantify the connection between bad respondents and the risk of drawing the wrong conclusion from survey data. The results were surprising:

- > Even a small proportion of bad respondents caused risk to increase exponentially.
- > As sample size increased, risk increased even more.
- > Eliminating only one type of bad respondent actually compounded the risk.

These findings are significant for all market research professionals but particularly for those conducting online research involving large sample sizes or using panels from multiple vendors. This paper summarizes the studies, our conclusions, and our recommendations for addressing the data-quality issue.

Data Quality: An Industry Challenge

As the use of online market research continues to gain momentum, research professionals are grappling with difficult questions about the quality of the data derived from online surveys:

- > Objectively, who actually participated in this study?
- > Was each survey taker for this study unique?
- > How engaged was each respondent throughout the survey?

MarketTools asserts that the market research industry—not just individual researchers—must directly address these issues to establish trust and confidence in the recommendations based on market research—and these questions need to be answered with hard data, not assumptions and suppositions.

Many of our colleagues and clients agree, but additional questions have arisen. Specifically, just how serious is the bad-respondent problem? What is the risk of making a wrong decision when only a small percentage of survey respondents is unreliable? When does it make sense to take corrective action, and what is the consequence of doing nothing?

Answering these questions with real data was the motivation behind recent studies undertaken by MarketTools—the first in the industry to quantify the connection between bad respondents and the risk of drawing bad conclusions.

Bad Respondents: A Closer Look

There is no formal, industry-standard definition of a bad respondent, nor is there consensus about the criteria for invalidating a survey respondent. For some research professionals, removing bad respondents simply means de-duping the panel. For others it means eliminating “speeders” and “straight-liners.” Still others focus on removing fraudulent respondents, such as those who attempt to game the system to collect incentives or for other motives.


MarketTools maintains that reducing the impact of poor-quality data requires a holistic approach to identifying and eliminating bad respondents. To that end, early in 2008 MarketTools introduced TrueSample,™ a technological solution that ensures data quality by verifying that each survey respondent is:

- > **Real.** Respondents must be who and where they say they are.
- > **Unique.** Respondents can never be allowed to enter a survey twice.
- > **Engaged.** Participants must provide honest, thoughtful responses.

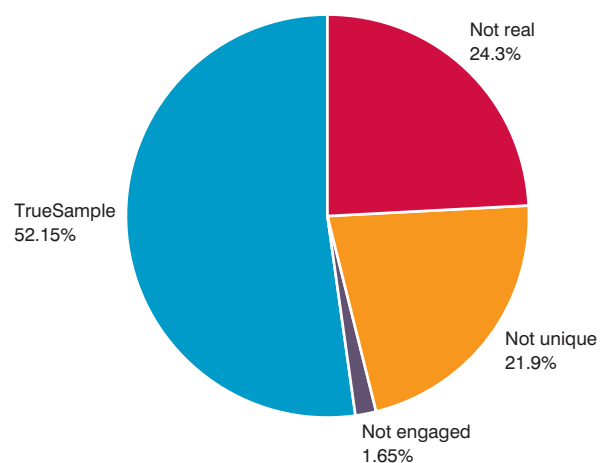
TrueSample brings the same real-time technologies that help prevent credit card fraud and identity theft to the world of online market research, enabling researchers to eliminate virtually all duplicate, fraudulent, or unengaged respondents from panel databases. Equally important, MarketTools applies those technologies not only to its own ZoomPanel™ but also to the panels of its Certified Partner network, upleveling the quality of survey data industrywide.

But is the holistic approach embodied by TrueSample truly necessary in all cases? What is the effect on the risk ratio if TrueSample technology is not applied, or if only one type of bad respondent is removed, or if nothing is done at all to remove bad respondents?

An average of 28.8% of panelists are rejected by TrueSample.



TrueSample rejects an average of 47.9% of panelists when any two panels are involved.



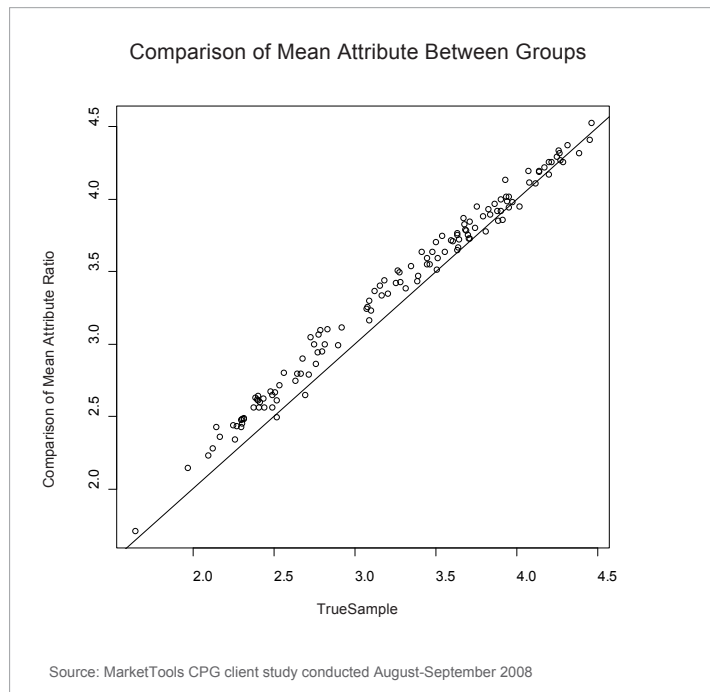
Source: MarketTools TrueSample processing August - October 08

The Experiment: What’s the Risk of Not Using TrueSample?

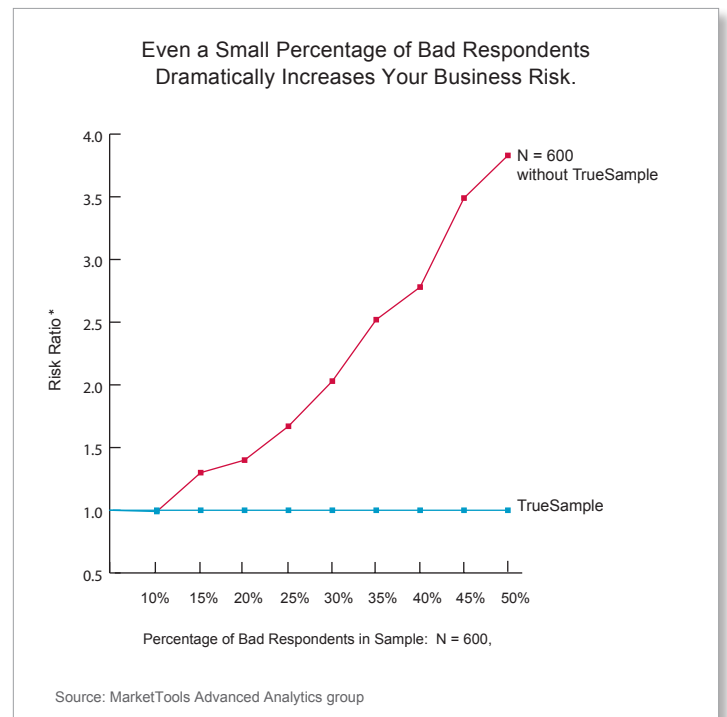
On behalf of a market-leading consumer packaged goods client, in the summer of 2008 MarketTools performed a comprehensive attitudinal study about consumer products used for breakfast. A total of 129 attribute ratings were used. We completed 304 interviews among respondents who had been certified with TrueSample, and we completed an additional 318 interviews with respondents who had been invalidated using TrueSample (identified as being duplicate or fraudulent).

One of the first questions we wanted to address was whether bad respondents actually answer survey questions differently from “good” respondents. After all, if there’s no real difference, there isn’t any increased risk of arriving at a wrong conclusion.

In a comparison of the mean attribute ratings between these groups, we found that the invalidated group gave consistently higher scores on the rating scales, as depicted in the diagram below.



When we mix in a small proportion of bad respondents—people who were invalidated by TrueSample—we find that the risk ratio increases exponentially. If a sample has 10% bad respondents, the increased risk is relatively small; but at 30% the risk is doubled, and at 40% the risk is nearly tripled.



100% higher risk: If your sample has 30% invalidated people, you have 2.03 times the risk of making the wrong decision—your risk is 100% higher.

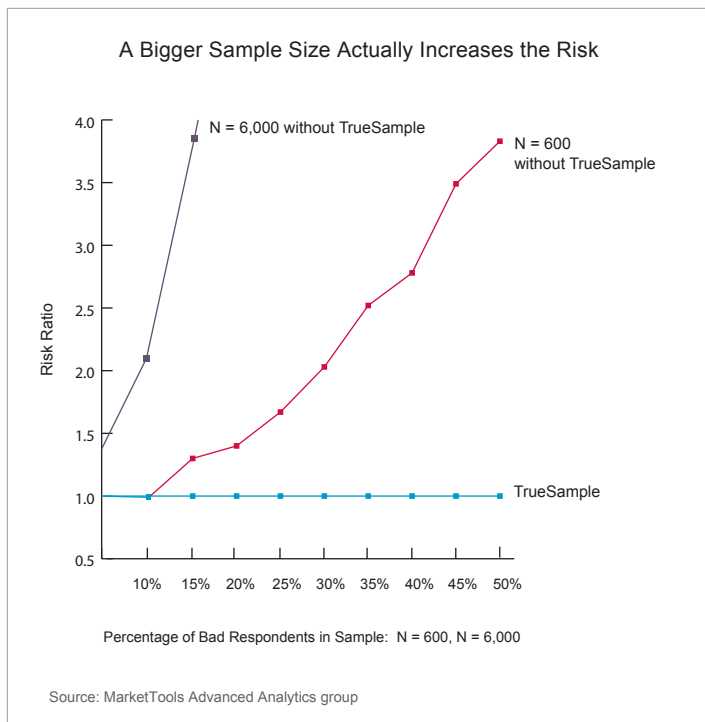
*A risk ratio of 1.0 indicates the same risk of error as using TrueSample, i.e. 5%. When the risk ratio doubles to 2.0 the probability of a wrong answer also doubles from 5% to 10%.

Respondents who were invalidated by TrueSample consistently gave higher ratings than the mean.

Bigger Sample Size, Much Greater Risk

One of the most surprising—and potentially alarming—findings of this study is that increasing the sample size actually increases the risk of a wrong decision. For example, if your sample has 30% invalidated people, you have more than twice the risk of making the wrong decision if your sample size is 600; and if your sample size of 6000 doubles the risk even if only 10% of your sample is invalidated.

This counterintuitive result can be explained by the fact that the small differences in how bad respondents answer questions is magnified as the sample size grows.



Higher sample sizes actually magnify the impact of small differences between “good” and “bad” respondents, dramatically increasing the risk of drawing the wrong conclusion.


Key questions to address at this point are: How many bad respondents does a typical sample contain? And how one can determine the percentage of bad respondents in a particular sample?

Industry estimates indicate that on average approximately 15% of sample is comprised of bad respondents. TrueSample typically eliminates about 20% of respondents, based on approximately 1600 studies to date.

Of course, sample quality varies widely across the industry, and MarketTools has seen a range from as little as 5% to as much as 64%. The percentage for any given sample depends on the vigilance of the panel health management practices; but for most companies, the percentage of bad respondents in the sample is sufficient to increase the risk of inaccurate conclusions.

B2B Sample Has High Risk Ratios as well

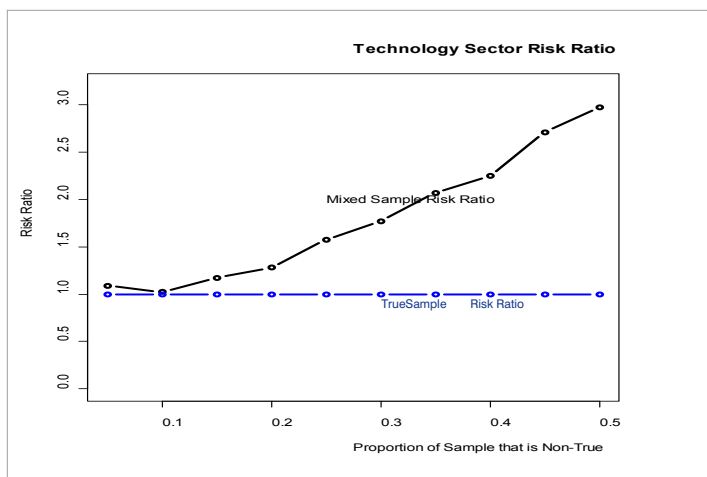
While the bulk of online research studies focus on the consumer, many companies are interested in targeting hard-to-reach (and heavily incented) B2B survey populations. We studied the impact of TrueSample via a study for a high-technology product, where the respondents were knowledgeable business users of the product. In this study, there were 30 attribute statements about the usage of different features and attributes of the high-technology product. The result: once again, the analysis showed significant upward bias in the ratings of the bad respondents—people who were invalidated by TrueSample.



B2B Respondents Also Show Upward Bias: Comparison of mean ratings

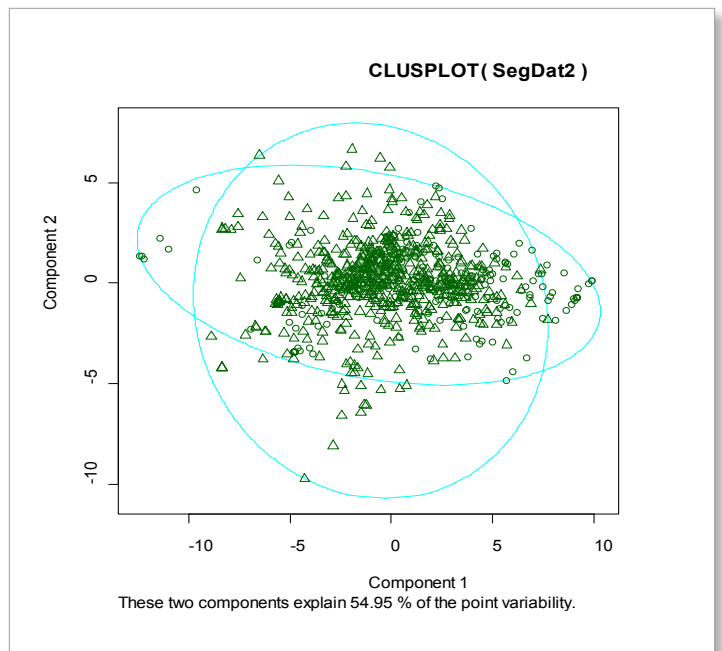
Risk Ratio Calculation for B2B Sample

When we calculate the Risk Ratio for this group of respondents, we get similar results to our consumer example—the risk of making a wrong decision goes up dramatically as the percent of bad respondents increases. It is important to realize that in respondent populations with highly specialized and compensated respondents, the proportion of “non-True” respondents can easily be higher than that observed in standard consumer samples.



Major Errors Are Possible

The impact of including these respondents with biased ratings has a huge effect if the purpose of the study is segmentation. In this case, using these 30 attributes, if we run a latent class segmentation algorithm on just the TrueSample respondents, the data does not support any subsegments, that is, the optimal number of segments is one. However, if the Non-TrueSample respondents are included in the task, we find that the latent class algorithm suggests two segments and most of the members of the second segment are “non-TrueSample” respondents.

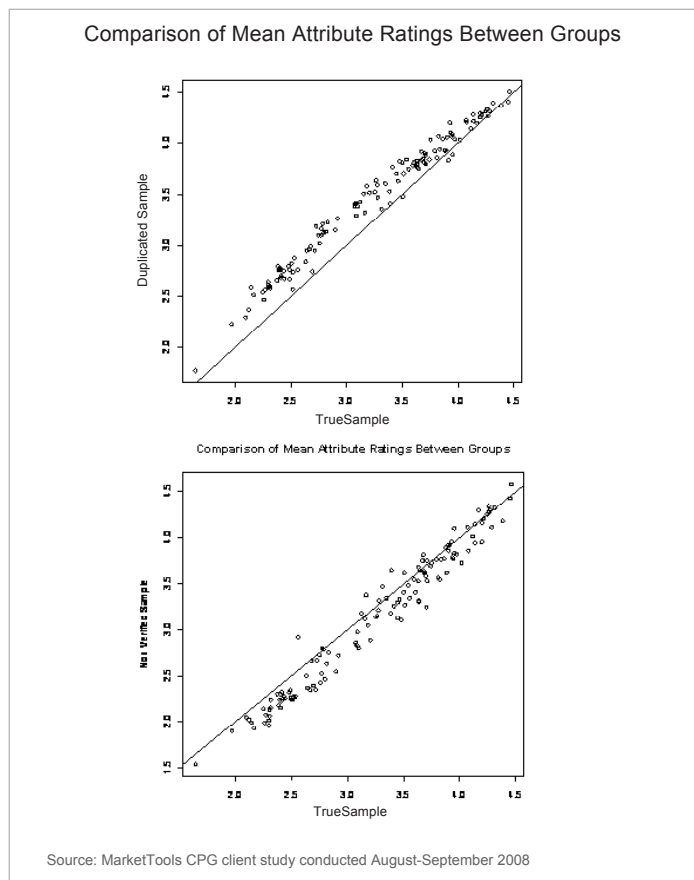


A Partial Solution is No Solution

Many companies attempt to mitigate the impact of bad respondents by focusing on one type of bad respondent at a time; that is, removing duplicates, then eradicating straight-liners, then going after those who are attempting to game the system to collect incentives.

Our studies show that addressing the issue in a piecemeal fashion actually exacerbates the problem.

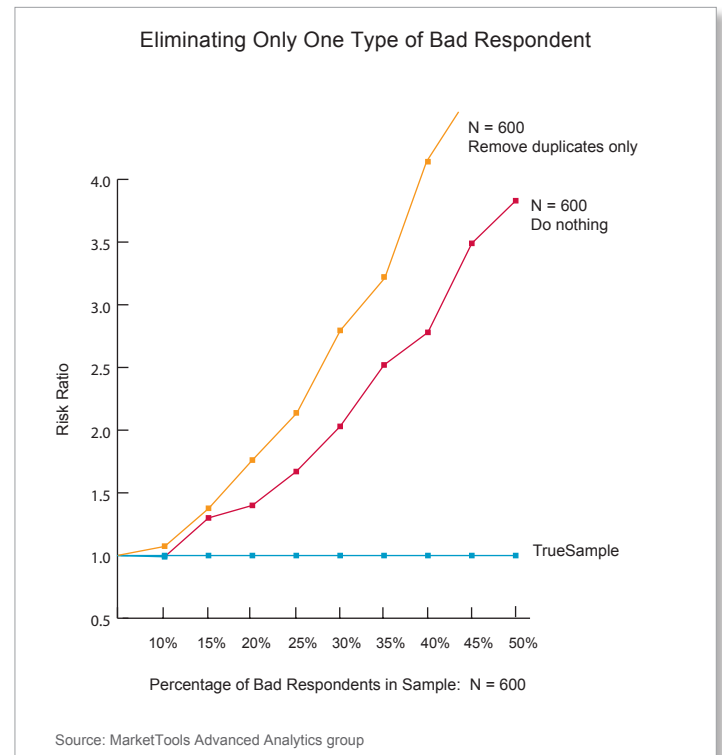
The two charts below show that two different groups eliminated by TrueSample both give bad answers—but in different ways. The first group, which consists of respondents identified by TrueSample as duplicates, consistently provided higher rating scores. The second group, identified as fraudulent by TrueSample, consistently provided lower rating scores. It is not relevant whether the ratings are higher or lower; the fact that the answers deviate from the norm is the significant finding.



The chart below illustrates what happens when only one type of bad respondent is removed.

With TrueSample, which removes all types of bad respondents, the overall risk ratio remains at an acceptable 5% (represented by the straight line at risk ratio 1.0). When no action is taken against bad respondents, the risk increases exponentially with the percentage of bad sample.

But when only one group of bad respondents is eliminated (such as removing only duplicates), the risk ratio increases at an even greater rate.



Eliminating only one type of bad respondent actually increases risk more than taking no action at all.

This finding raises a couple of questions:

1. Why would the response patterns differ between these two groups?

One explanation is that duplicates—people who have signed up multiple times—are likely to be people doing surveys to accumulate points or rewards. They therefore tend to give responses that are higher than normal on the hopes that a more positive set of responses will qualify them for further surveys or for free product trials.

Nonmatching people—whose identities cannot be confirmed—are more likely to be trying to hide their identity. Some may work in the market research field; others want to get an idea of what these online surveys look like. They are therefore more likely to give random answers, which, because of the human tendency toward agreement, will be lower on average than the answers supplied by TrueSample respondents.

2. Why does removing one group exacerbate the risk?

Because the groups have opposite tendencies, the mixture of the two pulls the average closer to the TrueSample average. Removing either group will pull the average score farther from the TrueSample average for the same proportion of bad respondents. If you could actually control the mix of the two types of bad respondents, you could approximate the TrueSample averages. Including both sets, however, will increase the variance on all your measures, as shown earlier in the section on, “The Experiment: What’s the Risk of Not Using TrueSample?” making your statistical tests less precise than they should be. Furthermore, if you believe that these groups are potentially giving bad information (as indicated by their differing response patterns), why would you want them in your study?

TrueSample: A Solution for Cross-panel Data Quality

TrueSample has proven effective in reducing the risk of deriving a wrong conclusion based on bad survey data from a given panel. Many clients, however, find that a single panel is not sufficient to meet their sampling needs, and they are concerned about the risks of duplication and/or fraudulent activity across multiple panels from multiple vendors.

TrueSample is the first industry solution that makes it possible to detect and remove bad respondents of all types across multiple panels. With TrueSample, researchers can objectively, measurably, and repeatably improve the quality of survey responses and online data sourced from a diverse range of vendors.

The TrueSample Certified Partner network now includes Greenfield, Luth, Toluna, and others. These companies have joined MarketTools in validating their panels using the TrueSample objective validation criteria—and in ensuring that our joint clients receive unique sample.

Technological Underpinnings of MarketTools TrueSample

TrueSample combines validation, digital fingerprinting, and de-duplicating technologies and applies them to ZoomPanel and the panels of all companies participating in the TrueSample Certification Program, ensuring that each respondent is “real, unique, and engaged.”

Real. MarketTools verifies the information that panelists provide against extensive databases built for specific industries, such as, consumer financial services, with objectively validated consumer demographics.

Unique. Panels from MarketTools and all companies participating in the TrueSample Certification Program are de-duplicated against one another. MarketTools goes a step further and uses digital fingerprinting to eliminate and blacklist fraudulent respondents across the TrueSample Certified Partner network.

Engaged. With our data validity technology, we can correlate survey-taking time and response patterns to identify fraudulent behavior and remove offenders from the data set and the panel. Further, with our SurveyScore™ technology, we can provide an objective measure of survey quality and engagement.

Conclusion

TrueSample significantly reduces the risk that poor data quality will translate to poor business decisions. By identifying and eliminating all sources of bad data, TrueSample mitigates both the obvious risks of using data from unreliable sample and the hidden risks of using large sample sizes to offset the effects of bad respondents. It provides a comprehensive solution to the data-quality challenge, and the net result is a measurable increase in confidence in the conclusions drawn from your research.

For the online market research industry, we believe that TrueSample represents a way forward—that it is the key to maintaining high-quality research and building trust in this dynamic, exciting era of accelerating market growth.

Michael Conklin

Chief Methodologist

Michael Conklin is a marketing research veteran and analytical methodology pioneer with over 20 years experience. Michael has deep knowledge in advanced analytics and is a Bayesian research modeling expert. In his capacity at MarketTools, Michael leads the Advanced Analytics teams and personally consults with clients on strategic research issues, providing sound analytical recommendations leading to deeper insights.

Prior to joining MarketTools, Michael served as Senior Vice President, Analytic Services for one of the top five global marketing research firms where he was instrumental in developing key analytical procedures for customer satisfaction, brand equity and product development. He is a frequent speaker at the American Marketing Association’s Annual Advanced Research Techniques Forum as well as at the American Statistical Association’s Joint Statistical Meetings conference. In addition, Michael has written numerous papers on research methodologies and applications for key marketing publications such as Marketing News and Marketing Research Magazine as well as numerous statistical journals.

About MarketTools

MarketTools is the leading technology and service provider of Customer Insight Management solutions for the world’s market leaders. Through a unique combination of best-in-class insight and feedback platforms, quality-assured global panels, and research innovation, MarketTools enables companies to better identify new opportunities, fuel greater product success, and build customer advocacy. As the first company to make online surveys widely available on the Web, MarketTools continues its market-leading position by providing the broadest range of powerful, accurate, and integrated customer insight technologies that empower companies to become the most customer-centric organizations in their industries. MarketTools’ premier portfolio of technology-based insight brands includes MarketTools.com,[™] Zoomerang,[™] CustomerSat,[™] TrueSample,[™] and ZoomPanel.[™]

MarketTools is a privately held company with corporate headquarters in San Francisco and European headquarters in London. For more information please visit www.markettools.com.



150 Spear Street
Suite 600
San Francisco, California 94105

main 415.957.2049
fax 415.957.2180
toll-free 866.499.3750

web www.markettools.com
e-mail sales@markettools.com