

BREWERY OPERATIONS BENCHMARKING SURVEY (BOBS)

The Brewers Association (BA) conducted a comprehensive survey of America's craft brewers to establish and share industry benchmarks and brewery operations data. This survey was developed in response to BA member requests for had data that they could use to measure their businesses and in turn make them more efficient and successful.

2014 survey questions were revised slightly from the 2012 and 2010 editions. Results were cleaned, compiled and aggregated by BA staff. The results provide a valuable set of tools specific to America's small and independent craft brewers.

Thank you to all who participated and may you find value in these results to help you grow with the industry. This is, indeed, an exciting time to be a part of the craft beer community!

Cheers,

A handwritten signature in black ink, appearing to read 'Bart Watson', is positioned above the printed name.

Bart Watson

Chief Economist

Background

Due to the sensitive nature of much of the BOBS information, the first goal of data compilation was to provide anonymity for all survey responses. This means that any statistics seen as overly revealing about a particular company or individual may be redacted or presented in a more limited form. This is particularly true for results related to salary, which reflect not only company, but also personal information.

While the need for anonymity was the first principle of compilation, the second was utility to members. Consequently, different statistics and different categories are used across different sections to provide the most useful information while preserving anonymity. The goals in shifting categories and statistics vary, but primarily focused on having meaningful sample sizes and statistics. For example, weighted means were used in areas where category averages may be as useful to brewers and allied trade suppliers as simple means (see the following section for more on the various statistics used and their meanings).

The majority of the data is separated into brewpub and production brewery categories, as well as by the size of annual production. As sample size increases, the data is presented in narrower production categories. For questions where the business model (brewpub vs production) was deemed less relevant, brewpub and production breweries are combined to improve the statistical power of the results.

Finally, the data has been extensively cleaned to ensure that responses reflect the spirit of the question being asked. At times responses may have been eliminated if they were deemed outside the bounds of reality and an easy data cleaning solution could not be found. Thanks to the support of many BA staffers in interpreting and cleaning survey responses, particularly Chuck Skyeck, Chris Swersey and Ryan Farrell. They deserve full credit for their work, although I assume final responsibility for any errors in the data presented.

Statistics/Terms Used in Results

Mean: The “average” value of a category. All the responses added up and divided by the total number of responses. Note that the mean is sensitive to “outliers” – values that stray far above or below the typical response.

Weighted Mean: Weighted means are akin to category averages (not an average of the category). These take into account the size of the breweries responded and may be particularly useful in differentiating breweries at the top of a production range from those at the bottom. See the example below for more.

Median: The middle response or the 50th percentile. The median is useful when you want a “typical” experience, without worrying about extreme outliers.

Example of mean vs weighted mean vs median:

Let’s use the following question with made up data. “What % of your volume is sold at the brewery?”
Let’s assume we got the following five responses:

Brewery A (50 barrels a year): 100%

Brewery B (100 barrels a year): 95%

Brewery C (500 barrels a year): 25%

Brewery D (1000 barrels a year): 30%

Brewery E (10,000 barrels a year): 3%

In this case each statistic would produce a very different result. The *mean* would be 50.6%, the *median* 30%, and the *weighted mean* 7.5%. These variations show the difference in each. The mean is pulled by the large values at the top, the weighted mean pulled heavily by the largest brewery in the category, and the median splits the values.

Mode: The most common response. A good statistic for the “what is the most common thing for a brewery to do?” question. Only useful when there are limited number of standard responses.

Min: The smallest value reported.

Max: The largest value reported.

Sample Size: The number of responses in a given category. Note that for many statistical purposes, the larger the sample size, the more confidence the observer should have about the tendency of the statistic to generalize to the larger population.

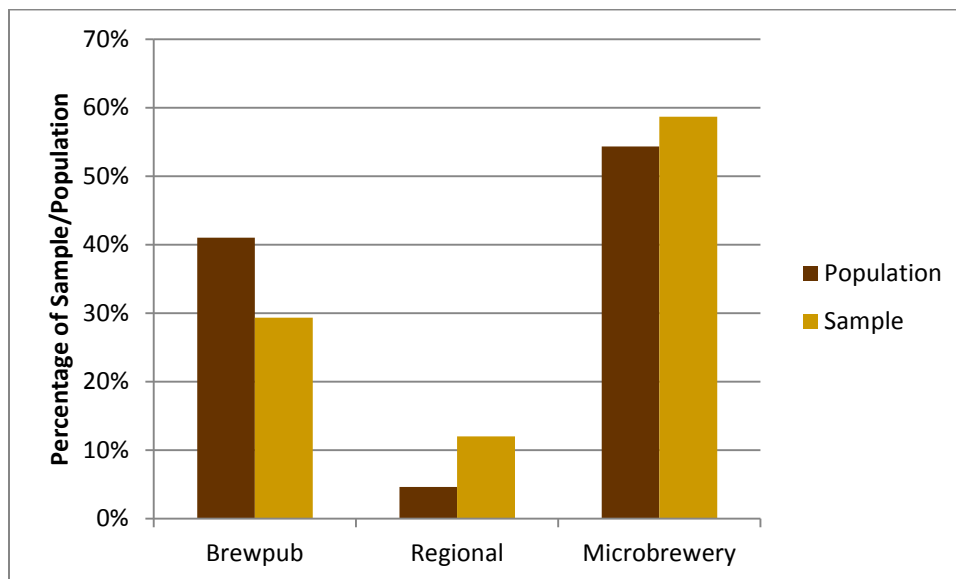
Percentile (inc.): A measure of a particular position in the distribution of responses. For example, Percentile 10 measures the answer that was in place 10 out of 100 (or 1 out of 10). Percentile 100 is the top value and Percentile 0 is the bottom. The (inc.) means that the percentile is inclusive, so it will include a value that falls at that cutoff. Percentiles are often used when given the full range of values (max and min) would reveal sensitive information.

Demographics of Respondents

One of the most important considerations of any non-randomized survey is that the sample be roughly representative (i.e. the respondents look generally like the larger population they are attempting to measure). Because brewers choose whether they want to fill it out or not, it is possible that BOBS might be vulnerable to a selection or “self-response” bias – i.e. that the people who fill out the survey are somehow different than breweries in general.

The BA received usable responses to all or part of the survey from 310 unique breweries (though the majority of questions have fewer total usable responses). This includes brewers from 46 states. No responses are included from North Dakota, Oklahoma, South Dakota, or Rhode Island.

The respondents were on average larger than the total brewery population, driven by an over-representation from regionals in answering the survey.



However, because responses are “stratified” (i.e. broken up into) groups based on business model and size, this is only a concern when results for all breweries are presented. Breaking results apart by business model and size largely eliminates worries that over-representation by regionals will skew the data.

Even with these caveats, all findings should be taken with a note of caution: given the possibility of self-report bias and low sample sizes for many questions, ***some results may not accurately reflect the whole industry or even sub-categories therein.*** Users are cautioned to use benchmarking data as one tool amongst others in formulating business plans and informing their business decisions.