Launching new tech? How do you make data-driven decisions without any sales data?

November 15 2018, by Liz Entman

Credit: CC0 Public Domain

When a tech company decides to launch a brand new, first-of-its-kind tech product, it can be hard to know how many you need to order. Order too few, and you may have to resort to more expensive manufacturing or shipping options to keep up with demand. Order too many, and you've just got a lot of wasted inventory on your shelves that you'll never sell at full price. Either way: Getting your product lifecycle forecast wrong can be an expensive mistake.

Kejia Hu, assistant professor of operations management at Owen Graduate School of Management, has developed a new method for forecasting the life cycle of new tech products that incorporates both historical sales data of predecessor products and business insights that is more accurate than current approaches—in some cases by a lot.

The research appears online in the *Journal of Manufacturing and Service Operations Management*. Her collaborators are Jason Acimovic at Penn State, Doug Thomas at University of Virginia, Jan A. Van Miegham at Northwestern, as well as Francisco Erize at Dell Inc.

Most new tech products aren't actually new—they're usually next-generation versions of things a company has made before. Estimating the life cycle of those products isn't too challenging because odds are good that Widget 5.0 is going to have a similar life cycle to versions 4.0 and 3.0, and so on. "The historical data from predecessors' sales information over their entire lifecycle will become a very powerful predictive data source for the next-generation demand," Hu said.

Tech products, which generally have intentionally short lifespans, see a sharp ramp-up of initial demand, followed by either a brief plateau or a

single high point, followed by a longer decline in interest as potential customers start waiting for the next-generation version of the product instead. If you plot that demand on a chart, it usually looks like a lopsided trapezoid or triangle. The actual numbers may differ from generation to generation, and the nuances of the curve will differ from product to product, but the general shape of the curve will likely be the same.

So how do you make data-driven decisions about a product so new that it has no data behind it? When that happens, companies usually rely on a combination of market research and the expertise of the product manager to build a forecast. That approach doesn't always work very well, however, so Hu and her co-investigators wanted to find a way to integrate hard data with those business insights to achieve a more accurate picture.

The solution, Hu proposed, is to use data from existing products with similar features or from similar categories as a kind of proxy for prior generations. For example, if a computer company is launching its first tablet, it doesn't have any other tablet data to look at. But it may have data on, say, a very lightweight laptop, a laptop with a touchscreen and a very inexpensive laptop close in price to the new tablet. This "cluster" of similar products forms the pool of data from which a product manager can draw to develop a curve.

"But beyond the data component, we also allow the top manager to overlay their perspective or their feelings for this new product," Hu said. "For example, if they think this is going to be a great product, they would probably do more promotional events, which will change some of the demand patterns along the way. So we also allow our forecast to be flexible enough to incorporate those business insights such as the planned launch time, promotional campaigns or planned sales events."

This blended approach can also be used to forecast demand for products that may still be too young to have much data behind them, like second-generation versions, and even to improve purely data-driven forecasts for well-established products with many generations of history to draw from.

Hu and her colleagues then tested their model using data from Dell and a small gaming hardware company called Turtle Beach. At Dell, Hu's model improved on Dell's forecast by an average of 3.4 percent for brand-new products, 9.2 percent for relatively young products, and 14 percent for established products, saving the company anywhere from $1.50 to $4.70 per unit of products. Meanwhile, at Turtle Beach, which is a small firm that doesn't necessarily have the resources to produce forecasts as sophisticated as Dell's, Hu's approach improved accuracy for established products by a stunning 73 percent.

Though Hu and her colleagues developed their model for tech products, she said the basic framework of their approach could translate to any new product designed to have a short life cycle, such as fast fashion, or even newer products with longer life cycles that don't have much sales data to draw on yet.

"The methodology is universal," she said. "But the need for this framework becomes more urgent when the product life cycle is short, so companies can really prepare for that one shot of sales."

More information: Kejia Hu et al. Finalist—2017 M&SOM Practice-Based Research Competition—Forecasting New Product Life Cycle Curves: Practical Approach and Empirical Analysis, *Manufacturing & Service Operations Management* (2018). DOI: 10.1287/msom.2017.0691

Provided by Vanderbilt University