



Predicting Inventory Value at Everything But the House's Online Auctions



EVERYTHING BUT THE HOUSE

CUSTOMER PROFILE:

Everything But the House is an online estate sale company that sells more than 80,000 unique items each month. Prior to deploying a model in the DataScience.com Platform that can estimate the value of its inventory, the company had no way of knowing the revenue it could expect to earn on a particular sale. Now, decision makers have access to financial forecasts that accurately predict monthly sales within cents.

HEADQUARTERS:

Cincinnati, Ohio

INDUSTRY:

Retail

WEBSITE:

www.ebth.com

EMPLOYEES:

1000+



THE CHALLENGE:

Predicting Revenue Across Hundreds of Auctions

Cincinnati, Ohio-based online estate sale company Everything But the House (EBTH)'s sole source of revenue is the commission it receives from more than 450 monthly auctions of items with a starting bid of just \$1. An auction can contain hundreds of unique items — EBTH sells an average of 80,000 privately-owned vintage pieces each month — and bidders have complete control over the final price. This combination of factors made forecasting revenue and understanding the performance of EBTH's online auction platform difficult.

To distinguish between auction performance and the inherent value of items being sold, as well as predict whether revenue targets would be met, senior leadership needed a data science-driven solution. Using the DataScience.com Platform, which is now the hub for all data science work at EBTH, Data Analyst Lavneet Sidhu built a model that predicts item value by averaging closing bids on 10 similar items. Similar items are identified using data from their description, categories, and title. Two months after being deployed in the DataScience.com Platform, the model was able to predict the total dollar amount of monthly sales with over 90% accuracy.



THE SOLUTION:

Scheduling Code Runs and Sharing Results in a Data Science Platform

EBTH's entire data science team is using the DataScience.com Platform to build, run, and deploy models, as well as communicate their results to leadership — no small feat at a 1,000-person company that serves 20 markets across the country. At present, the item value model, which Sidhu built in a Jupyter Notebook on the platform, is used for short-term financial forecasting. The results are primarily shared with stakeholders in a Tableau Workbook.

“Our senior leadership monitors these price predictions closely,” says Business Intelligence Architect Jonathan Kyrlach. “This assists us in establishing a baseline for what we expect for the outcome of a particular auction, but it also gives us the ability to say, ‘This didn’t sell for what we expected. Why is that?’ Maybe the sale wasn’t promoted correctly or it wasn’t accessible enough to potential bidders.”

Ultimately, Kyrlach hopes the model will be used to identify which sellers the company should add to its marketplace by assessing the projected value of their inventory.

In addition to leveraging the platform’s compatibility with Tableau, the data science team is also taking advantage of the Scheduled Runs feature, which allows users to run a script at regular or custom intervals using the hardware size and environment of their choice. The item value model is run twice daily and updated predictions are automatically stored in a database.

“It takes a significant amount of time for the model to process items based on similarity,” Sidhu notes. “The ability to schedule jobs has been tremendously helpful — now I can let the model run in the background and update the database directly.” Prior to implementing the DataScience.com Platform, Sidhu and his teammates were only able to run analyses such as these on an ad hoc basis.

“ Until we deployed a price prediction model in the DataScience.com Platform, we could only wait and assess post-mortem auction performance. It has provided us with a capability that we just didn't have before. It's extremely valuable.”

JONATHAN KYRLACH

Business Intelligence Architect

THE OPPORTUNITY:

Scaling Data Science Across EBTH

The predictions being generated by Sidhu’s model are informing EBTH’s dynamic sorting model, which is deployed as an API in the DataScience.com Platform. Previously, employees were manually updating the order of items on auction category web pages based on instinct, a process that took hours and was causing excessive overtime. The company’s product team decided a major change was necessary: Now, the website pings the API deployed in the platform when a sale goes live to retrieve its order number, which is partially determined by its projected value.

EBTH’s data team works closely together to identify opportunities like these. As the team grows, collaborative features like a company-wide view of work done in the platform, the ability to invite teammates to projects, and one-click report publishing from notebooks, HTML, and Markdown files will make it easy for teams to break down silos and start reusing, reproducing, and sharing code.

“Until we deployed a price prediction model in the DataScience.com Platform, we could only wait and assess post-mortem auction performance,” Kyrlach says. “It has provided us with a capability that we just didn't have before. It's extremely valuable.”