

How to Manage Scheduling Software Fairly

by Ethan Bernstein, Saravanan Kesavan, and Bradley Staats

FROM THE SEPTEMBER 2014 ISSUE

Starbucks workers recently scored a point against the machine. After a lengthy *New York Times* story, the company decided to adjust some of their controversial scheduling practices, eliminating “clopening” – when workers are required to close at night and re-open in the morning – and requiring at least a week’s notice of upcoming schedules.

In this case, “the machine” refers to a real machine: the highly sophisticated automated software Starbucks uses to schedule its 130,000 baristas, sometimes giving them less than a few day’s notice about their schedules in order to “optimize” its workforce.

Starbucks is just one of many companies using this type of technology, and it’s not hard to understand why. Until recently, determining who works when involved store managers manually slotting each employee into shifts on paper. Automation not only frees store managers to focus on customers, but can take into account much more data than a person can remember – historical customer patterns, weather, experience at neighborhood stores – so workers spend less time either with nothing to do or being completely overwhelmed with long lines and unhappy customers.

As the argument goes, that’s good for workers, who don’t want to be bored or overwhelmed, and it’s good for retailers, whose biggest cost and revenue driver is typically labor.

Our collective research has also shown that retailers really do struggle with scheduling. In a study of 41 stores of a women’s apparel retailer, for example, Saravanan Kesavan and his co-authors found that all of the stores were understaffed significantly during the peak periods of the day, while they were significantly overstaffed during the rest of the hours. The authors estimated that the retailer was losing about 9% of sales and 7% of profits due to this mismatch.

So why not implement just-in-time, software-driven staffing across the board? The problem – and one that Starbucks was forced to face first-hand – is that while scheduling software may seem “like magic,” as one of the major software vendors in the *Times* article put it, it’s actually not. Starbucks’ experience is common among a number of retailers who have taken their passion for engineering and optimizing schedules too far. As soon as a computer is scheduling your people at 15-minute increments to match the peaks and valleys of customer demand, employees’ desire to live a normal, predictable life becomes a barrier to profitability.

Three additional realities get in the way:

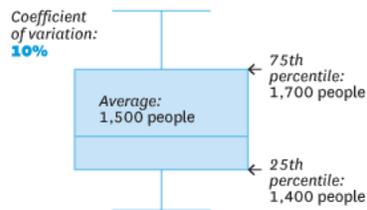
Perfect forecasts don’t exist. To produce an optimal labor schedule, the scheduling software must forecast customer patterns accurately—if you want to schedule labor at 15-minute increments, you must also understand demand at 15-minute increments. The dirty little secret is that even the most advanced scheduling software, incorporating every bell and whistle, tends to be wrong at least as often as it is right when the time intervals are short.

For example, the charts below show the variation in customer arrival pattern for a retail store on Saturdays, and then a breakdown of the same store between noon and 1 p.m. on Saturdays. The coefficient of variation, a measure of how variable the traffic is, increases from 10% to 31% as we go from daily to hourly data. This implies that it would be lot harder to predict hourly traffic compared to daily traffic, and it’s guaranteed to be lot more erroneous when predictions are generated at the 15-minute interval.

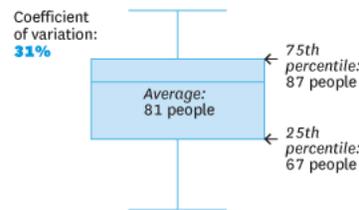
CUSTOMER ARRIVAL PATTERNS TO A RETAIL STORE

When measuring traffic over a shorter period of time, there is a substantial increase in variation.

SATURDAY (ALL DAY) STORE TRAFFIC



SATURDAY (NOON TO 1PM) STORE TRAFFIC



SOURCE DATA FROM A MEDIUM-SIZED APPAREL RETAILER, ANALYZED BY SARAVANAN KESAVAN.

HBR.ORG

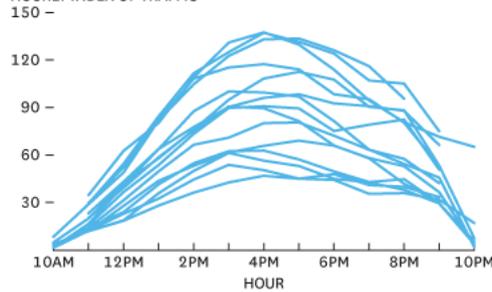
Because of this, good scheduling software tends to serve the “normal” customers well (those buying cappuccinos on the way to work every morning), but it may be at least as important to serve the abnormal ones (like a person who buys 20 lattes for a single meeting).

Tracking everything is unreliable. The response to unpredictable customer behavior has been, “well, let’s just track everything more carefully.” But tracking doesn’t always yield better predictions. For one, there’s inherent variability, as shown in the arrival pattern of customers to 35 stores of a retail chain over a period of a day. The differences between the top and bottom lines - in other words, the differences between two stores - don’t offer much in the way of insights.

HOURLY VARIATION IN STORE TRAFFIC

A retail chain’s customer arrival pattern based on normalized data from 35 stores.

HOURLY INDEX OF TRAFFIC



SOURCE DATA FROM A MEDIUM-SIZED APPAREL RETAILER, ANALYZED BY SARAVANAN KESAVAN.

HBR.ORG

Increased tracking can also create distortions in the data and unforeseen employee reactions, as Ethan Bernstein has found in his research. At a factory in Southern China, for example, executives thought that watching workers would help managers increase productivity and lower costs. It turns out that the opposite was true: Employees were more likely to be innovative when they weren’t being monitored, and production slowed when all eyes were on them.

A lot of flexibility isn’t necessarily a good thing. Saravanan Kesavan, Brad Staats, and their co-author have shown that temporary and part-time workers can help improve sales and profitability up to a point. For example, store profits were found to increase when the number of part-time and temporary workers was increased from zero to 4-5 for every 10 full-time workers. But beyond that point, any further increase in the number of those workers decreased store profits as the intangible costs of lower motivation and greater coordination dominated the benefits of scheduling workers to meet unpredictable demand.

So what should retailers do? To start, they can learn a lot from the considerable variation in how these systems have been implemented across companies. We see them as falling somewhere in the spectrum of purely creating *value for management* and purely creating *value for workers*.

The systems designed solely to create value for management fail to take into account so many unpredictable (and human) variables that they often result in public failures. In 2011, over 4,000 workers from Macy’s threatened to go on strike, in part because employees felt that the management was pushing for an online scheduling system that would make their schedules unpredictable. Walmart has often been accused by working mothers that the unpredictable schedules and low wages hurt their lives. The story of Jannette Navarro, the Starbucks worker profiled in the *Times*, absolutely rings true.

At the same time, strictly human-centric approaches can be problematic: If store relies on a remote scheduler, that's a person who isn't spending time with customers or employees. Like the scheduler Ms. Navarro had to beg in order to get 40 hours, he or she may also exert unfair power over workers.

Some systems, however, are implemented to create value for workers while simultaneously taking advantage of software's benefits. If a store has a scheduling system that is accessible to, and modifiable by a team of corporate managers, store managers, *and* workers, it can balance human needs, customer needs, and company needs in a fair, transparent, and more informed way. These implementations can make ordinary workers, even part-time workers, better at managing themselves.

In other words, the best kinds of scheduling systems involve both managers and software, not for the purpose of more tightly controlling workers, but to inform them on how optimal schedules stack up against predicted forecasts. For example, what if store-level employees could edit the schedules produced by the machine, but were held accountable for the ultimate effectiveness of them?

This is the exact approach taken by Belk, the largest family owned and operated department store in the United States. Before their tool was implemented, scheduling was performed by store managers and schedulers who balanced profits and worker needs to create a "fair" schedule that worked for everyone – incorporating preferences and an equal amount of weekend work. These types of nuances, with lots of variation, were too complicated for any workforce software to take into account.

So when Belk implemented their new tool, employees saw its failure to create fair schedules as "bugs" in the software.

But unlike other retailers who take an iron hand to push compliance with a new system, this retailer let the team at the store "edit" the system to "fix" the "bugs" – essentially, they allowed workers and their supervisors to ensure they had the days or hours off they needed, more than a week in advance, by overriding the system.

At the same time, a central workforce team at corporate was tasked with analyzing a sampling of edits to understand their reasons and benefits. Some edits were, of course productive; others involved resistance to change or misunderstandings and miscommunications. Belk then worked with its store managers through weekly meetings to encourage compliance in areas where the scheduling system made sense, and at the same time provided feedback to the scheduling company to update its software where the schedules did not make sense.

Belk now revises about 50% of its scheduling based on this new approach, a healthy balance between the efficiency you get from a machine and the intelligence you get from human intervention. And the company reported a 2% increase in gross profits several months after implementing the override system.

Ultimately, the success of scheduling systems depends on whether they serve as tools for or against the workers. In many ways, data-driven scheduling software is attractive to retailers because it gives them unprecedented transparency. But the ultimate success of these systems depends on this same transparency being available to employees as well. When management takes enabling its workers seriously – when these tools become an experiment in worker learning rather than top-down compliance – results can far exceed even the most magical predictions that scheduling software initially promised.

A version of this article appeared in the December 2014 issue of HBR.

Ethan Bernstein is an assistant professor of business administration in the organizational behavior unit at Harvard Business School and the author of "The Transparency Paradox" (*Administrative Science Quarterly*, June 2012). Twitter: @ethanbernstein

Saravanan Kesavan is an associate professor of operations at UNC Kenan-Flagler Business School.

Bradley Staats is an associate professor at the University of North Carolina.

This article is about MANAGING PEOPLE

 FOLLOW THIS TOPIC

Related Topics: TECHNOLOGY | RETAIL & CONSUMER GOODS

Comments

Leave a Comment

POST

0 COMMENTS

 [JOIN THE CONVERSATION](#)

POSTING GUIDELINES

We hope the conversations that take place on HBR.org will be energetic, constructive, and thought-provoking. To comment, readers must sign in or register. And to ensure the quality of the discussion, our moderating team will review all comments and may edit them for clarity, length, and relevance. Comments that are overly promotional, mean-spirited, or off-topic may be deleted per the moderators' judgment. All postings become the property of Harvard Business Publishing.