This new inventory management tool reaches 94% perfect stock allocation

Product stock in the fashion industry is highly perishable. Fashion items quickly lose value as trends shift, and seasonality plays a major role in demand. What's more, today's customers can choose from an increasingly wide range of apparel – and they're taking full advantage of that variety, becoming more selective than ever before.

Unlike in other industries, where it's possible to use a target stock level (TSL) model to calculate maximum demand, the highly uncertain nature of fashion demand requires inventory managers for fashion retailers to operate on much thinner safety stock margins, to avoid being left with a large stock of unsellable inventory at the end of an item's life cycle.

But such a thin margin also comes with an increase in risk. No retailer wants their available inventory to fall below market demand – especially at the peak of an item's popularity, when they stand to generate the greatest profit.

These challenges led the team at Evo to approach the problem of stock levels in a qualitatively new way, by developing a machine learning model that learns from a wide range of market signals, and forecasts demand for each item at each point of sale.

In fashion retail, traditional TSL models fall short.

By their very nature, TSL models work best with durable goods like electronics and canned foods, which can be overstocked and redistributed throughout the supply chain over an extended period.

The fashion retail industry, on the other hand, lacks the certainties that make such a TSL calculation effective. This sector's strong seasonality and unpredictability require a new kind of model: one that accurately predicts demand for each item in each store, and regularly updates itself to incorporate new market factors into its analysis.

Fashion retail demands a new approach to inventory optimization.

In order to make accurate predictions, a TSL model needs to start with an accurate estimate of supply chain scheduling. Without this estimate, the model can't predict how often to make reorders, or what the lead time on those reorders should be.

An effective TSL model also needs a thorough forecast of the demand curve, which is often impractical for fashion retail – and it needs regular updates to capture the current value of the items whose stock levels it's calculating.

For all these reasons, Evo set out to replace traditional TSL with a dynamic tool for stock level optimization. We started by asking the question, "How do we optimize stock levels if we don't know our demand distribution beforehand?"

The more we analyzed this question, the more we realized there's no single answer. Even when we start with a clear understanding of the initial allocation and delivery time, we still need to incorporate a variety of client-specific business rules, location-related variables, and sales variances between different items, sizes, times of day, and so on.

Our answer was to stop focusing on answering any single question. Instead, our model asks a series of six different questions, then incorporates all the answers into a single prediction of the ideal stock level for a particular item at a particular store at a particular time.

Let's illustrate this process with a quick example:

Measuring the performance of six different replenishment approaches

Say we start with 100 stores, which collectively have the potential to sell 1,500 items (average demand of 15 items per store). To keep things simple, we'll perform this analysis for just one item, at just one size.

See what happened when Evo set out to replace traditional #TSL with a dynamic tool for #fashion #inventory optimization. https://bit.ly/2tOy5ZG Tweet This

We'll assume we already know the demand distribution among the stores, which enables us to assign a 20-percent average variability for each store. We'll also assume we're working with an in-season product selling at full price, with no discounts or promotions.

Keeping those parameters in mind, let's take a closer look at the analysis. Each of our model's six approaches focuses on optimizing for a particular driver of inventory allocation:

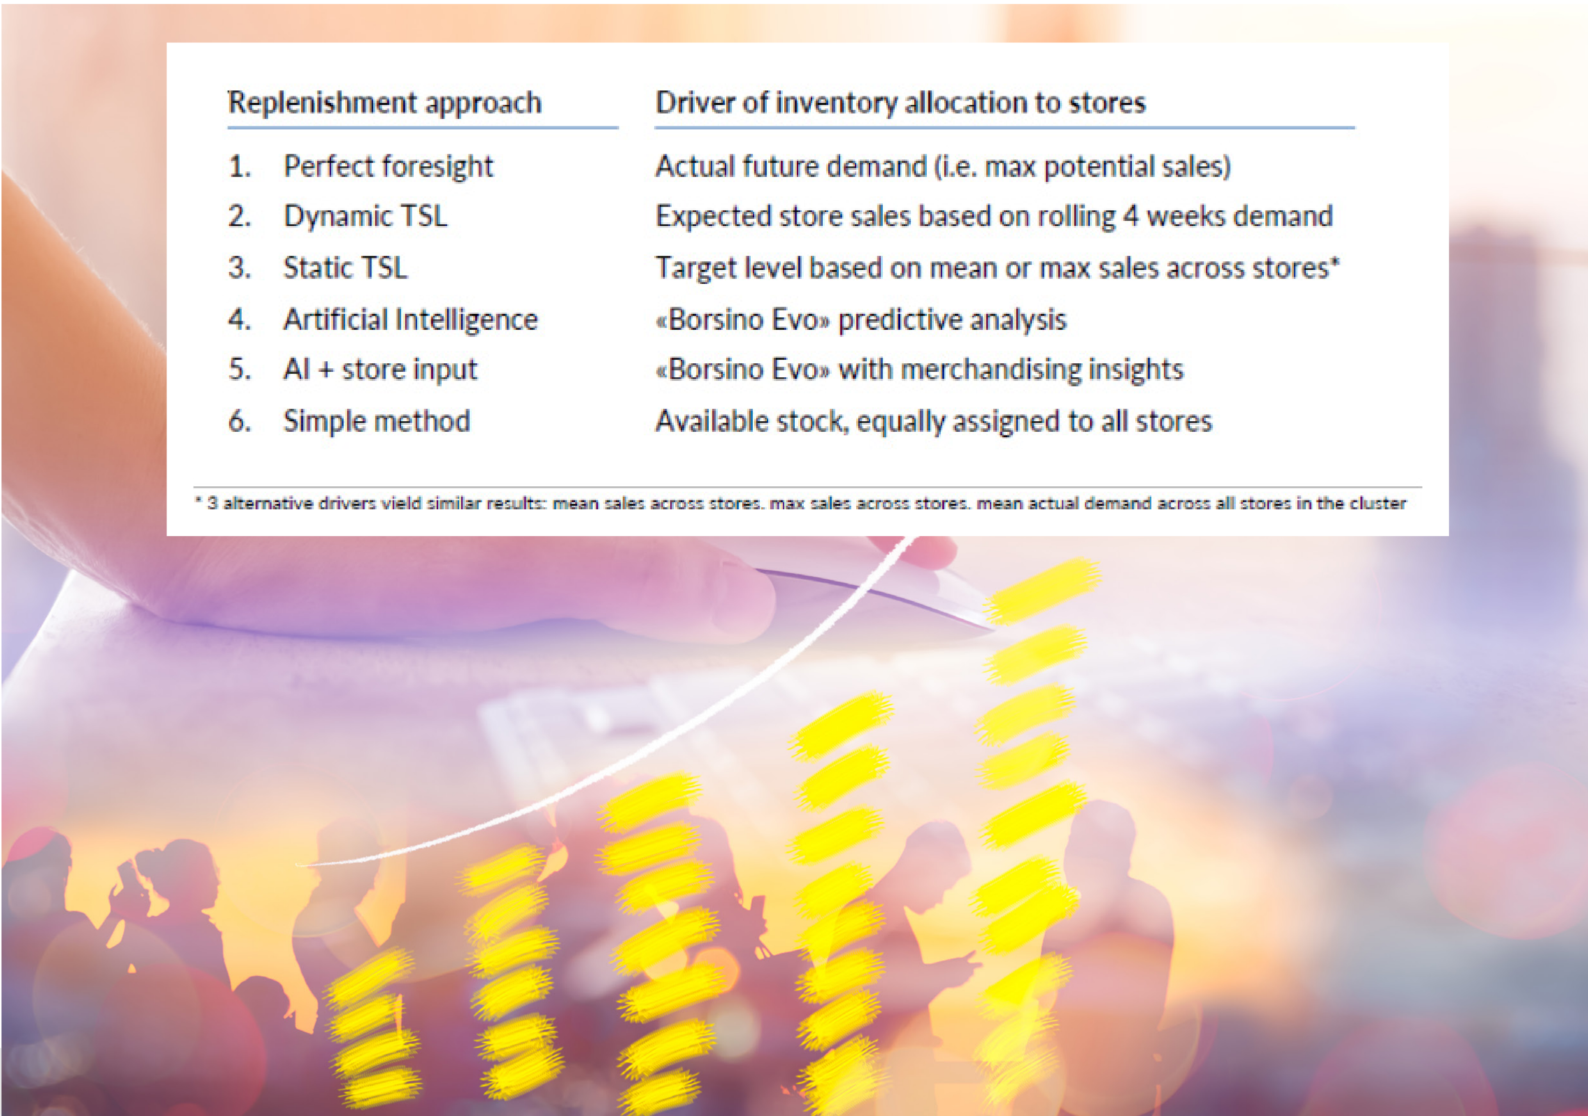

Replenishment approach	Driver of inventory allocation to stores
1. Perfect foresight	Actual future demand (i.e. max potential sales)
2. Dynamic TSL	Expected store sales based on rolling 4 weeks demand
3. Static TSL	Target level based on mean or max sales across stores*
4. Artificial Intelligence	«Borsino Evo» predictive analysis
5. AI + store input	«Borsino Evo» with merchandising insights
6. Simple method	Available stock, equally assigned to all stores

* 3 alternative drivers yield similar results: mean sales across stores, max sales across stores, mean actual demand across all stores in the cluster

By integrating the results of these six replenishment approaches, our machine learning model was able to generate a 9-to-25-percent gain in sales over traditional TSL methods.

The results improve even further when store managers edit their own allocation proposals. In fact, when our AI and store managers work together, our model's forecast achieves 94 percent of theoretically perfect sales performance:

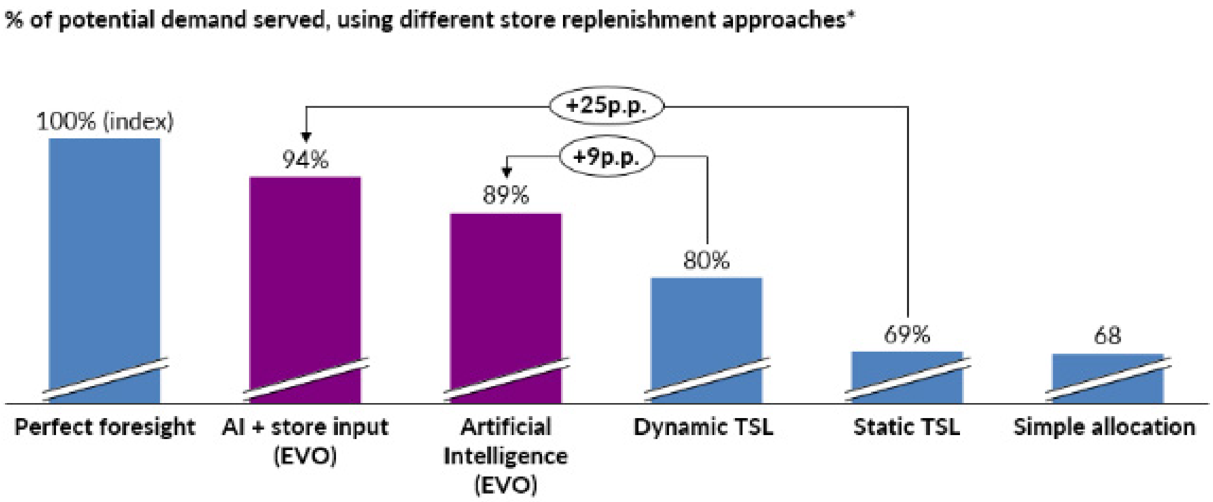

Breaking down the results of our stock optimization analysis

The figure below shows the the percentage of pieces sold, depending on the capability of the inventory to fulfil the total demand at that time:

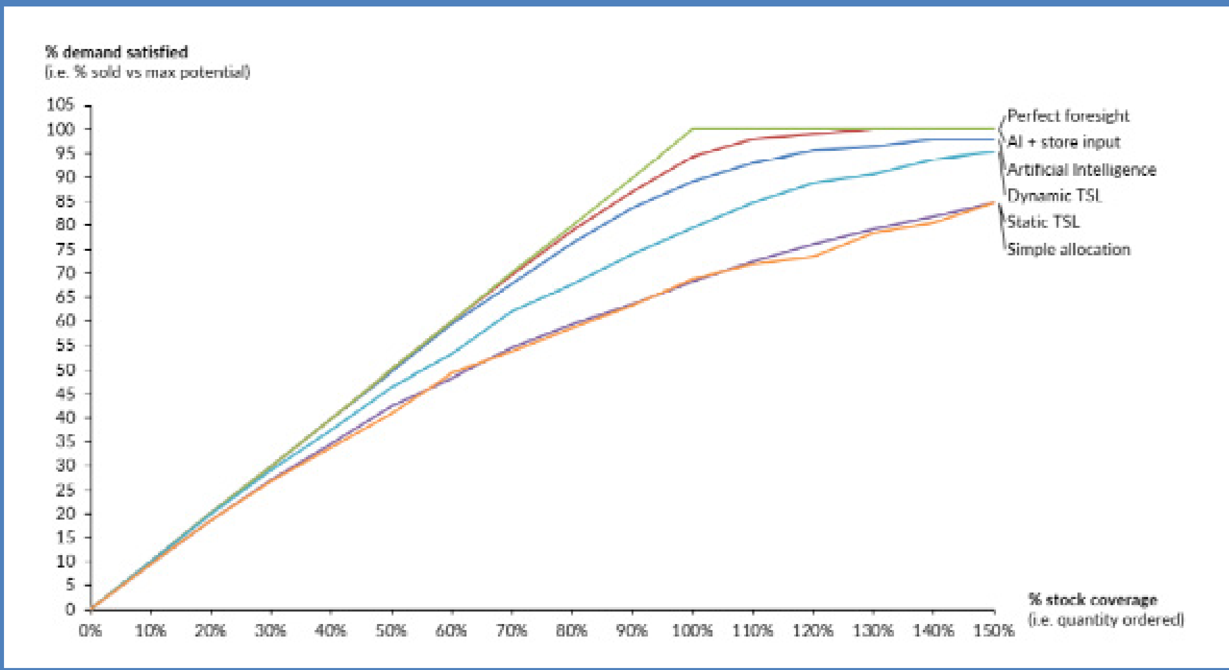

And the figure below shows the quantities remaining at the end of the period in question:

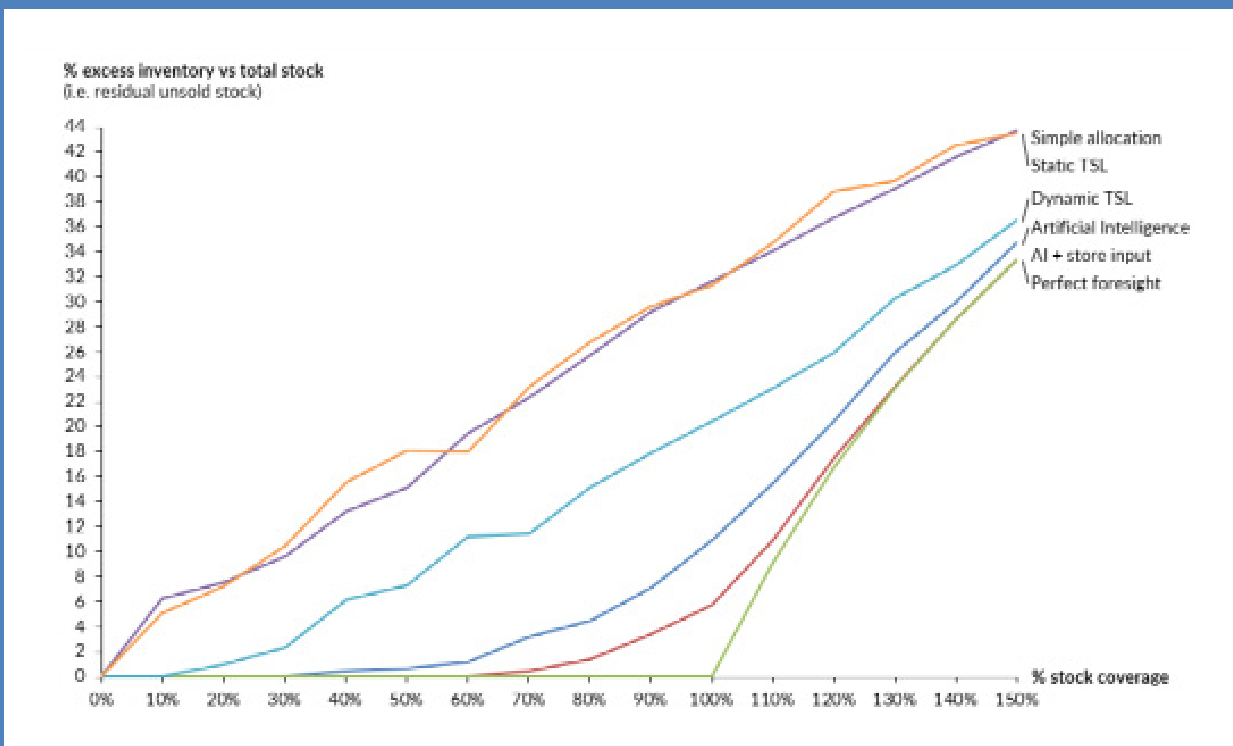

If the total availability is less (or more) than what's needed at a given time, the Dynamic TSL detaches from the benchmark curve of total foresight result at 10 percent of the total stock coverage, leading to stockout issues. This departure from ideal behaviour occurs at 60 p.p. (with human input) and 30 p.p. (without human input) – whereas it is almost immediate when static replenishment policies are implemented.

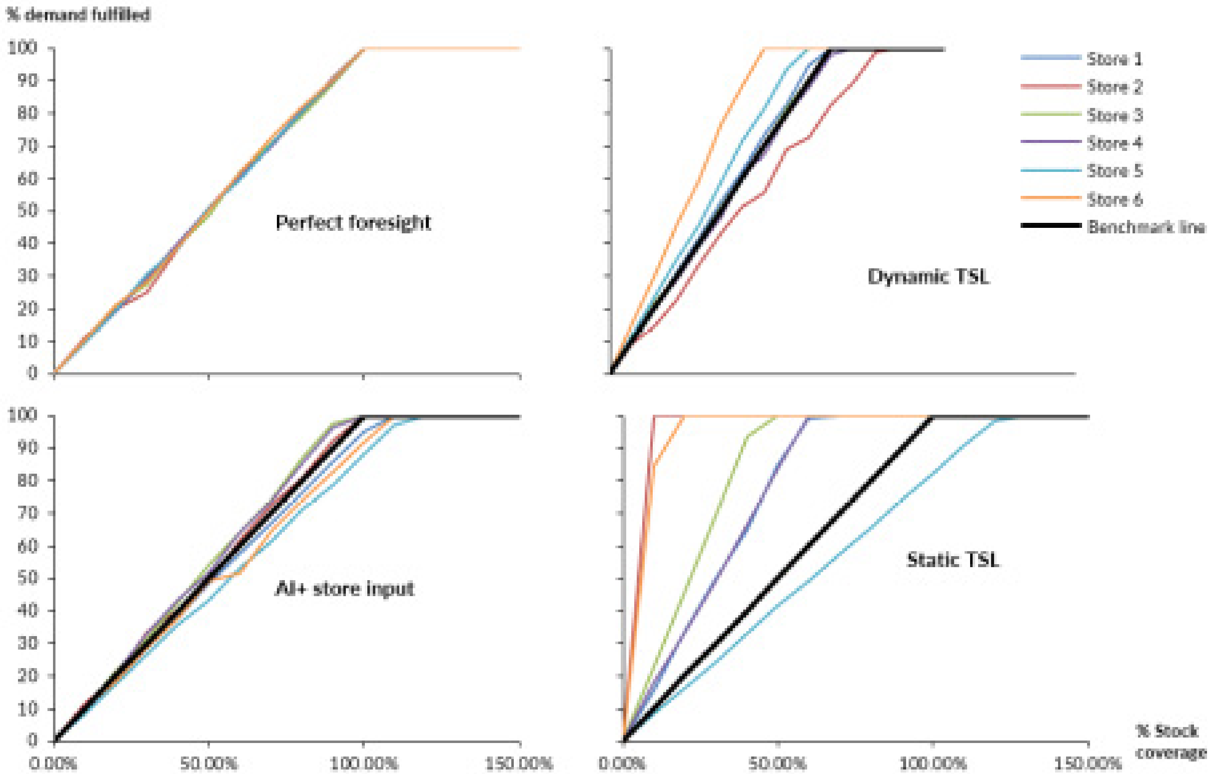

The lower the accuracy of the forecasting, the more the trend of each sales curve spreads and moves away from the bisector – as shown in the figure below, where static TSL leads to an overall loss of about 63 percent: One key learning from this analysis was that static inventory allocation is almost never a good thing. It's rarely able to meet demand in some locations, while other locations end up with far too much stock.

Second, we learned that a simple forecasting method (such as traditional TSL) isn't enough to obtain the best profit from the inventory on hand. The only effective way to generate optimal inventory forecasts is to regularly feed new data into the model, so it learns and optimizes its own calculations.

And third, we learned that AI on its own can't measure up to the performance of AI and human experts working in tandem. When staff members can input their own business rules, specify their own replenishment outcomes, and grade the results, the AI's performance comes very close to a "perfect foresight" stock allocation generated in retrospect, resulting in highly optimized stock levels for every item in every store.