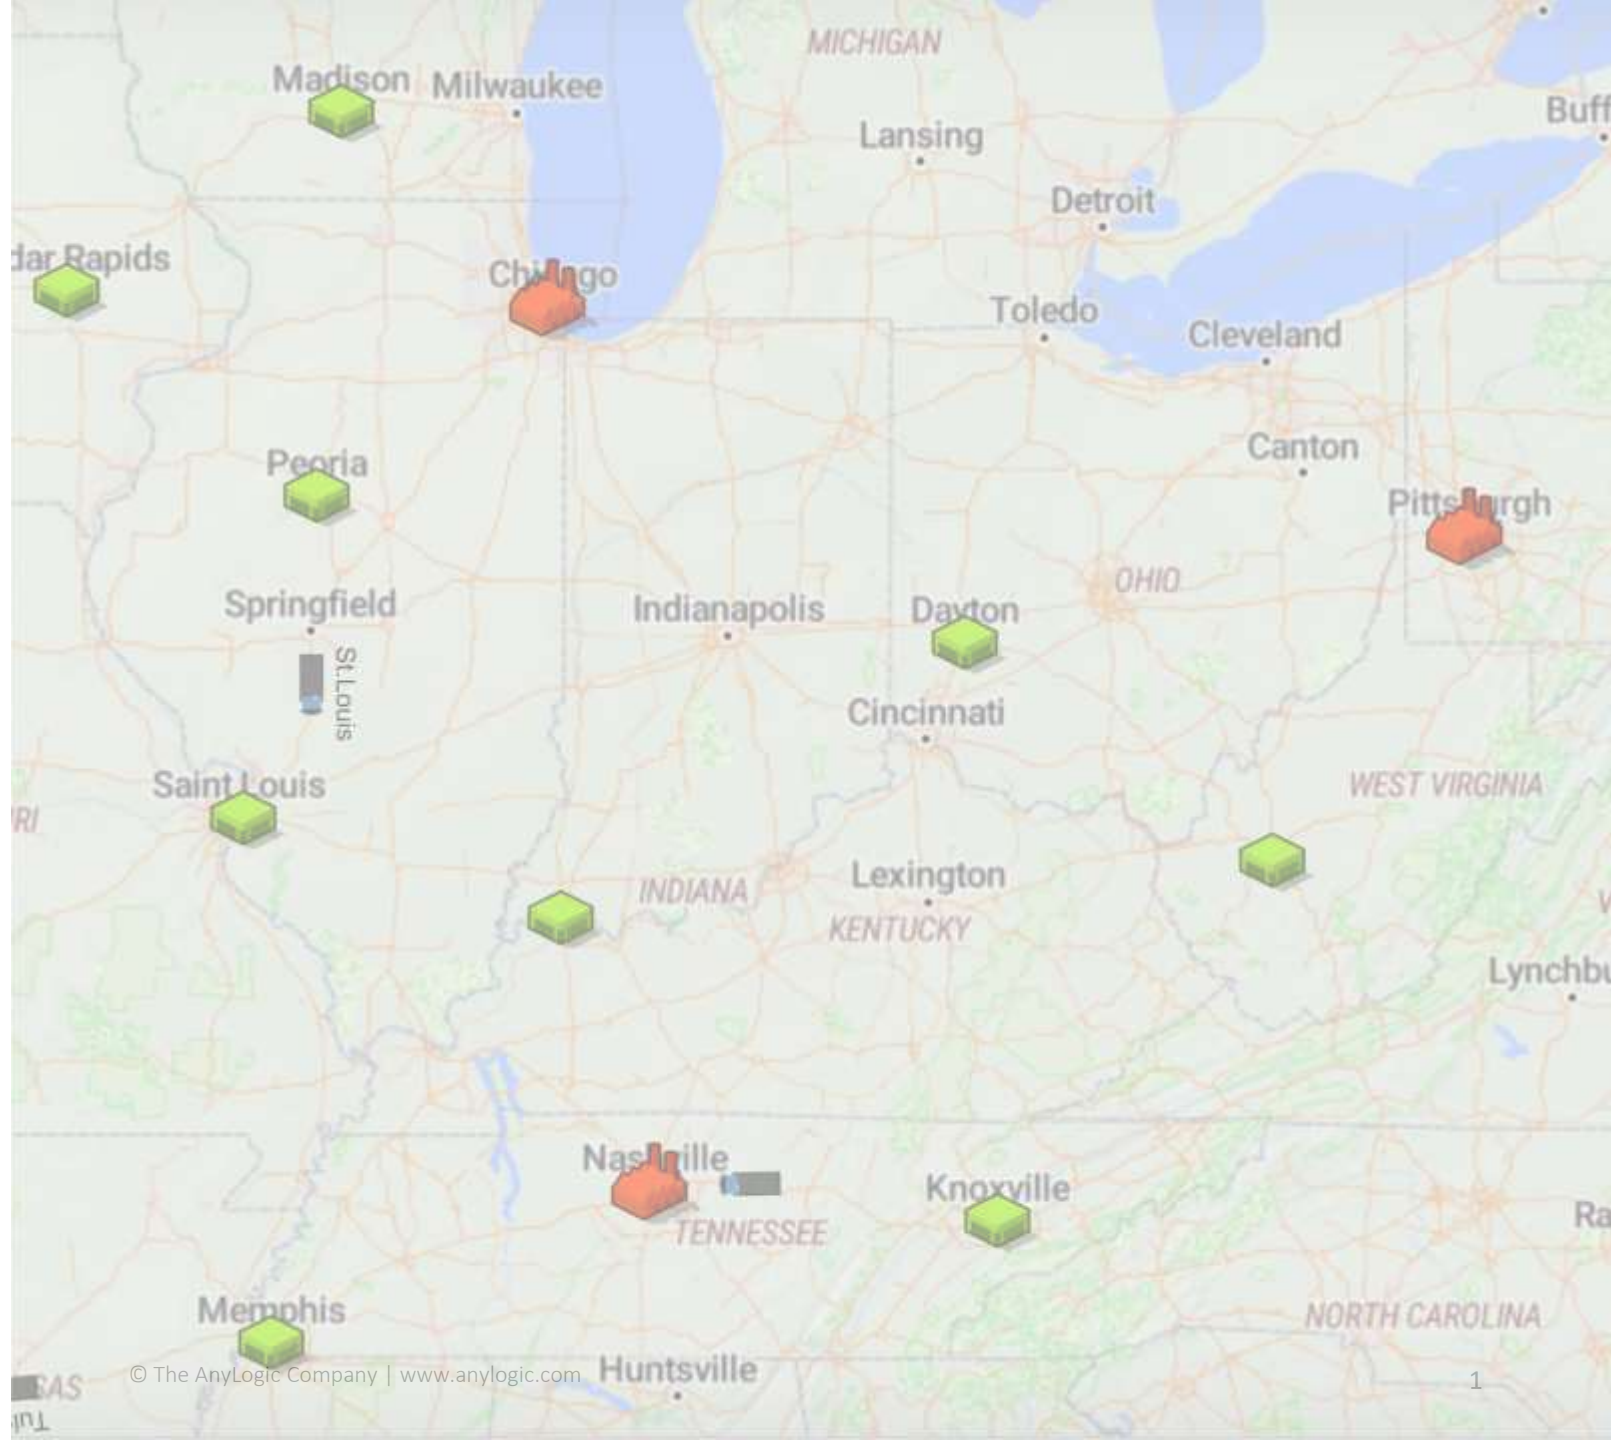


# Product Delivery

with Bonsai

-- an AnyLogic model



# Overview

- A multi-method model depicting a single-echelon supply chain – from manufacturing centers to distributors – serving a single product
- A “self-configurable” model: Agents are populated on startup based on the values in the built-in database
  - They are placed on a GIS map with trucks used to transport the completed orders traveling along real-world routes

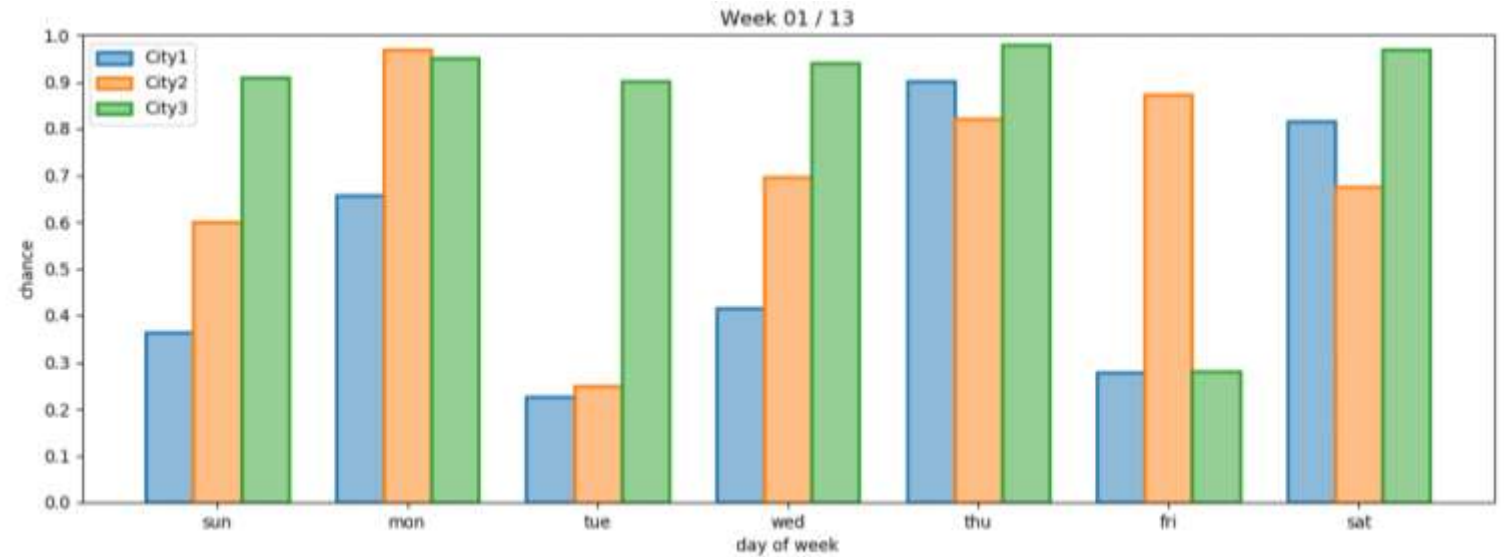


# Overview


- 3 manufacturing centers create the continuous product (e.g., liquid) and service incoming orders
- 15 distributors periodically send order requests to the nearest center
  - Each has a unique weekly schedule that varies on a macro-level over a fixed period



Sample chances to order over time

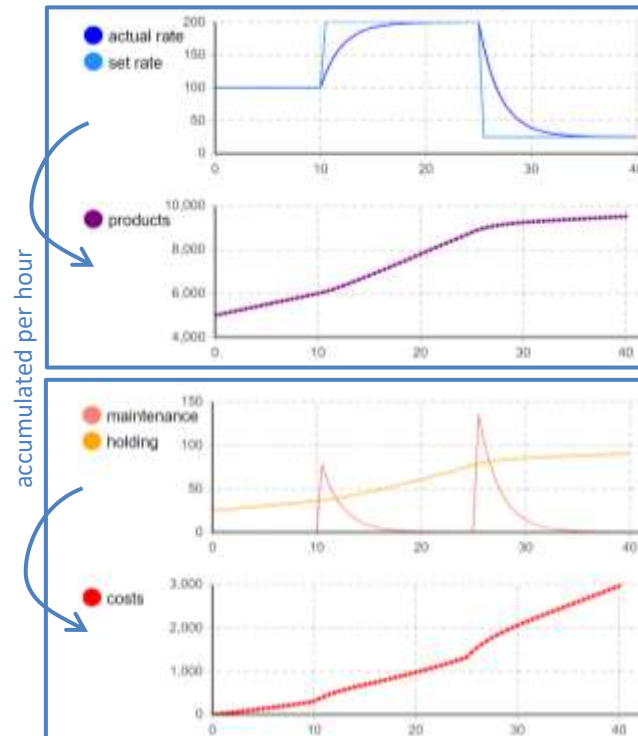



Sample order size distributions



# The Problem + Goal

- Each manufacturing center controls its assigned production rate
  - Machines require ramp-up time; modeled as a two-hour exponential delay
- Only one order can be serviced at a time due to an intensive manufacturing process results
  - If an order arrives and there is not enough product on hand, it – and any other subsequently arriving orders – must wait
- Two sources of costs are considered:
  1. Holding costs: The product is increasingly costly to store on-site
  2. Maintenance costs: For machine upkeep; typically nominal but more costly when changing the production rate
- Ideal conditions:
  - Keep enough product on hand to service orders rapidly, while keeping costs low



Sample depiction of products and costs (and their breakdown) from two rate updates over a 40-hour period.

For simplicity of showing this, no orders were accepted during this time.

# Current (Baseline) Solutions


## 1. Static rates

- A fixed approach in which each center has its production rate set to a constant value for the entirety of the simulation
- Provided as a worst-case baseline




## 2. Inventory Policy

- A linear approach, setting the production rate based on the current number of products on-hand



## 3. Heuristic

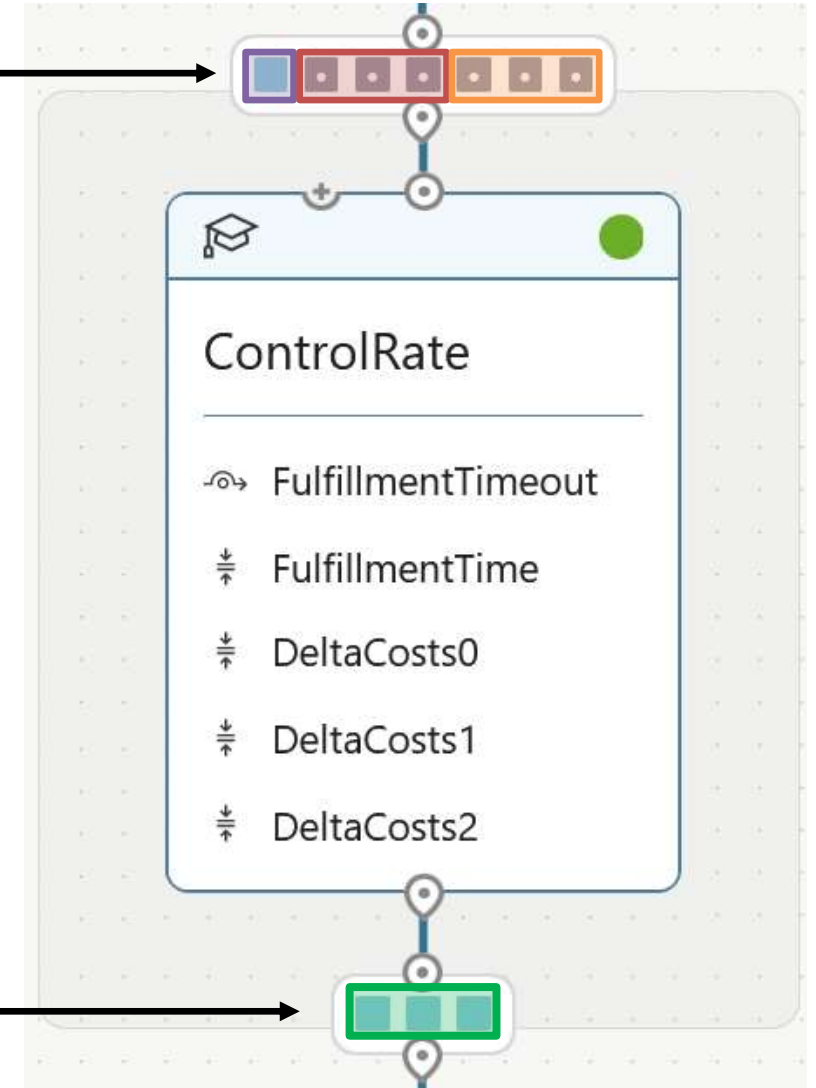
- A custom heuristic that seeks to hold a safety stock for one day ahead
- For planning, it uses the historical average for the given day
- Attempts to compensate if not on track for the current day's average amount



- These simple methods were chosen for their ability to be easily understood and later tweaked/maintained by humans, rather than their robustness for achieving the system's goals.
- A non-linear autonomous brain, able to easily handle the nonlinear configuration of this problem, is theorized to perform vastly superior.

# Brain Design / Development

- Episode duration: 13 weeks (default length of one demand cycle)
- Observable state:
  - – Current day of the week
  - – How many products on-hand at each center
  - – The cumulative amount of queueing products at each center
- Action – taken once per simulated day:
  - – Production rate for each of the three manufacturing centers
- Goals:
  - Keep fulfillment time below 24 hours; never allow it to exceed 48 hours
  - Keep costs per day below \$200
- Scenarios:
  1. Low variability in order patterns (limited range of RNG seeds)
  2. High variability in order patterns (full range of RNG seeds)



# Results

- A custom dashboard was created in AnyLogic to compare the baseline methods and any number of brains, in parallel
  - Each instance uses the same RNG seed such that the order patterns are the same
- Results from 1 year (4 order cycles) can be seen below; mean values are on the right of each graph
- The brain was able to capture the average weekly patterns and macro-level cycles to keep both costs and fulfillment time low!

