Pricing Algorithms, Nursing Homes, and Covid-19

Ben Tengelsen

Your task is to build a pricing algorithm

- Your clients are nursing homes – they need substitute nurses for specific shifts
- You employ a pool of nurses (10's of thousands), but they choose when/where they work
  - Facility agrees to pay a flat $X/hr
  - You fill the shift for $Y and you keep (or lose) the difference
- Requests can come many weeks in advance, or just 30 minutes
- Prices need to be set at the time the shift is requested (semi-dynamic)
- Lots of fun surprises - censored outcomes, diverse market conditions, opinionated executives

Where do you start?

Problem: Today's "prediction" data is tomorrow's "training" data

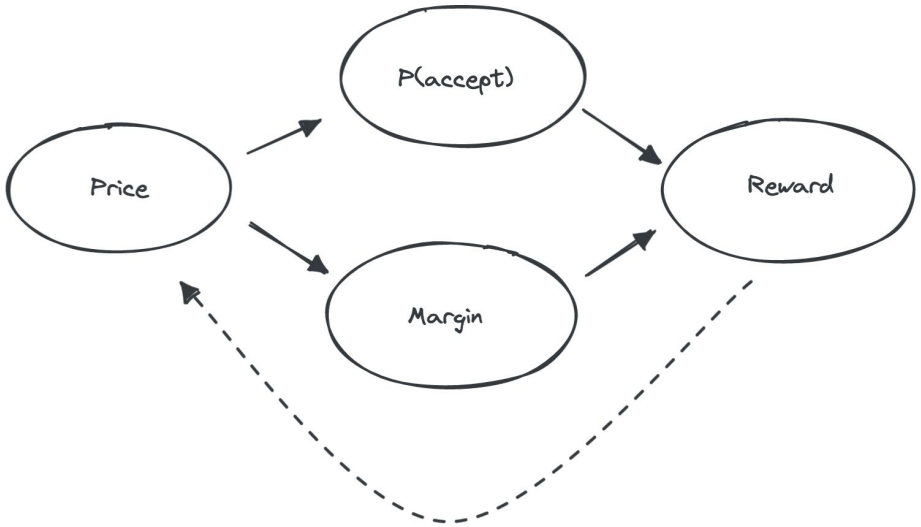

Bonus Challenges

- Continuous treatment space
- long-tailed purchase data
- Weird human behaviors (anchoring, habits)
- Short shelf life
- Today's prediction data is tomorrow's training data

Things are getting tricky!

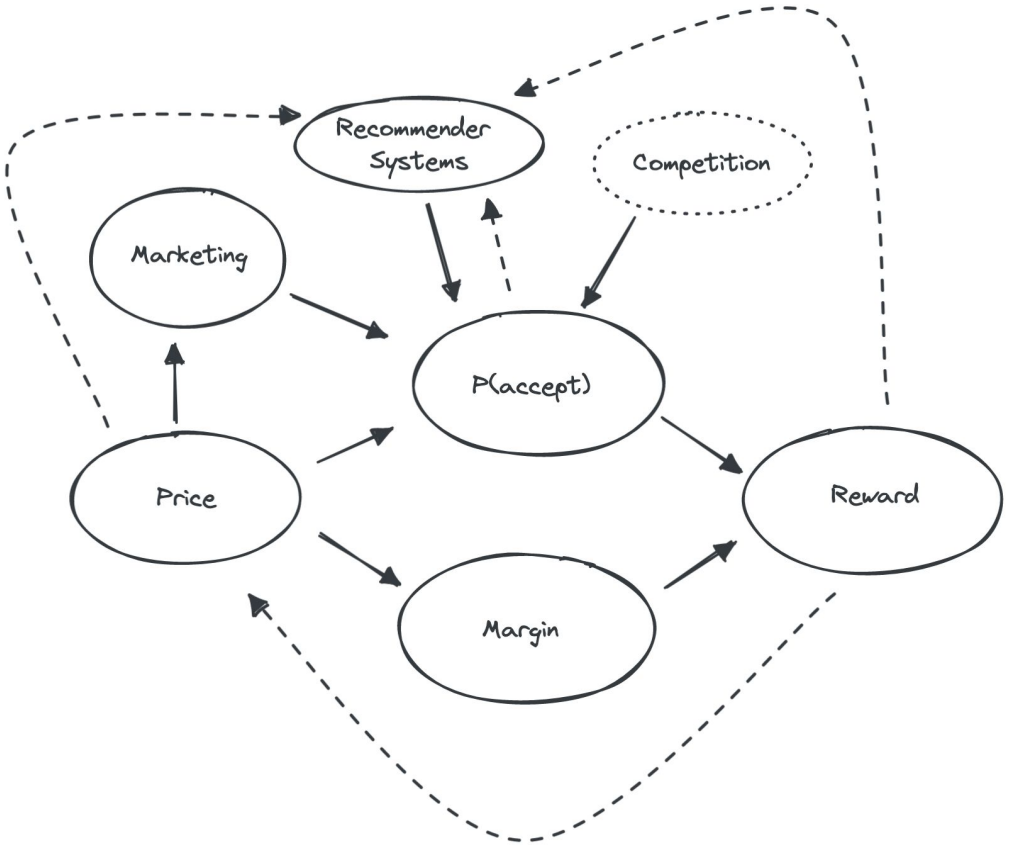

Bonus Challenges

- Continuous treatment space
- long-tailed purchase data
- Weird human behaviors (anchoring, habits)
- Short shelf life
- Today's prediction data is tomorrow's training data
- You probably don't observe
  - competition
  - specific marketing outcomes

Or just run some experiments!

Optimal pricing is about good ~~models~~ data

One experiment is not enough

- To set optimal prices all the time, we need good data anytime we train our model

Some popular ways to systematically generate quality data

- Random subset
- Bandits

For better or worse, we chose bandits

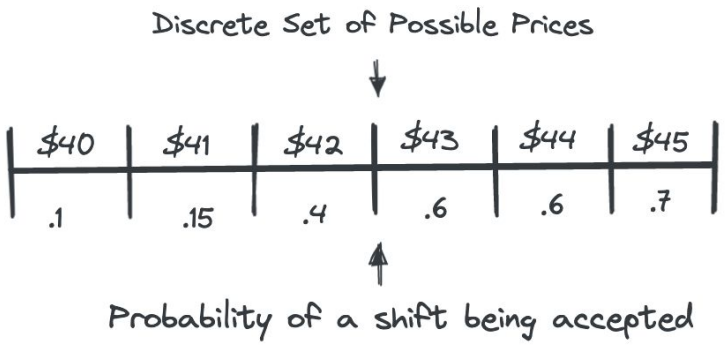

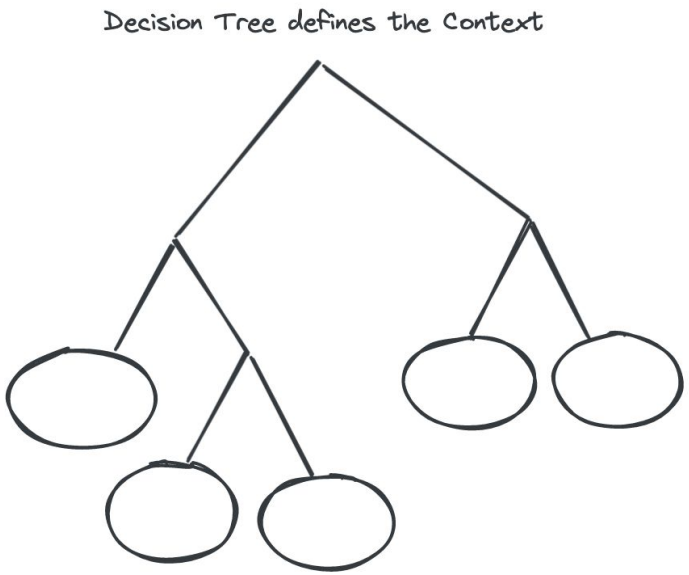

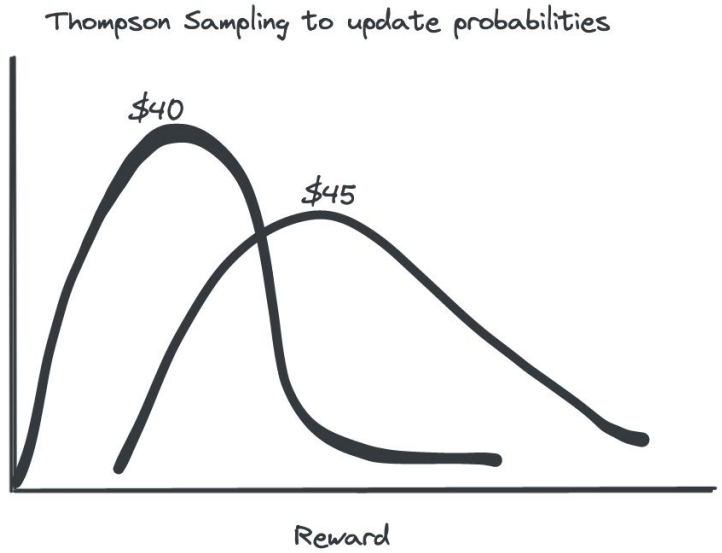

This is what it looks like (real chart)

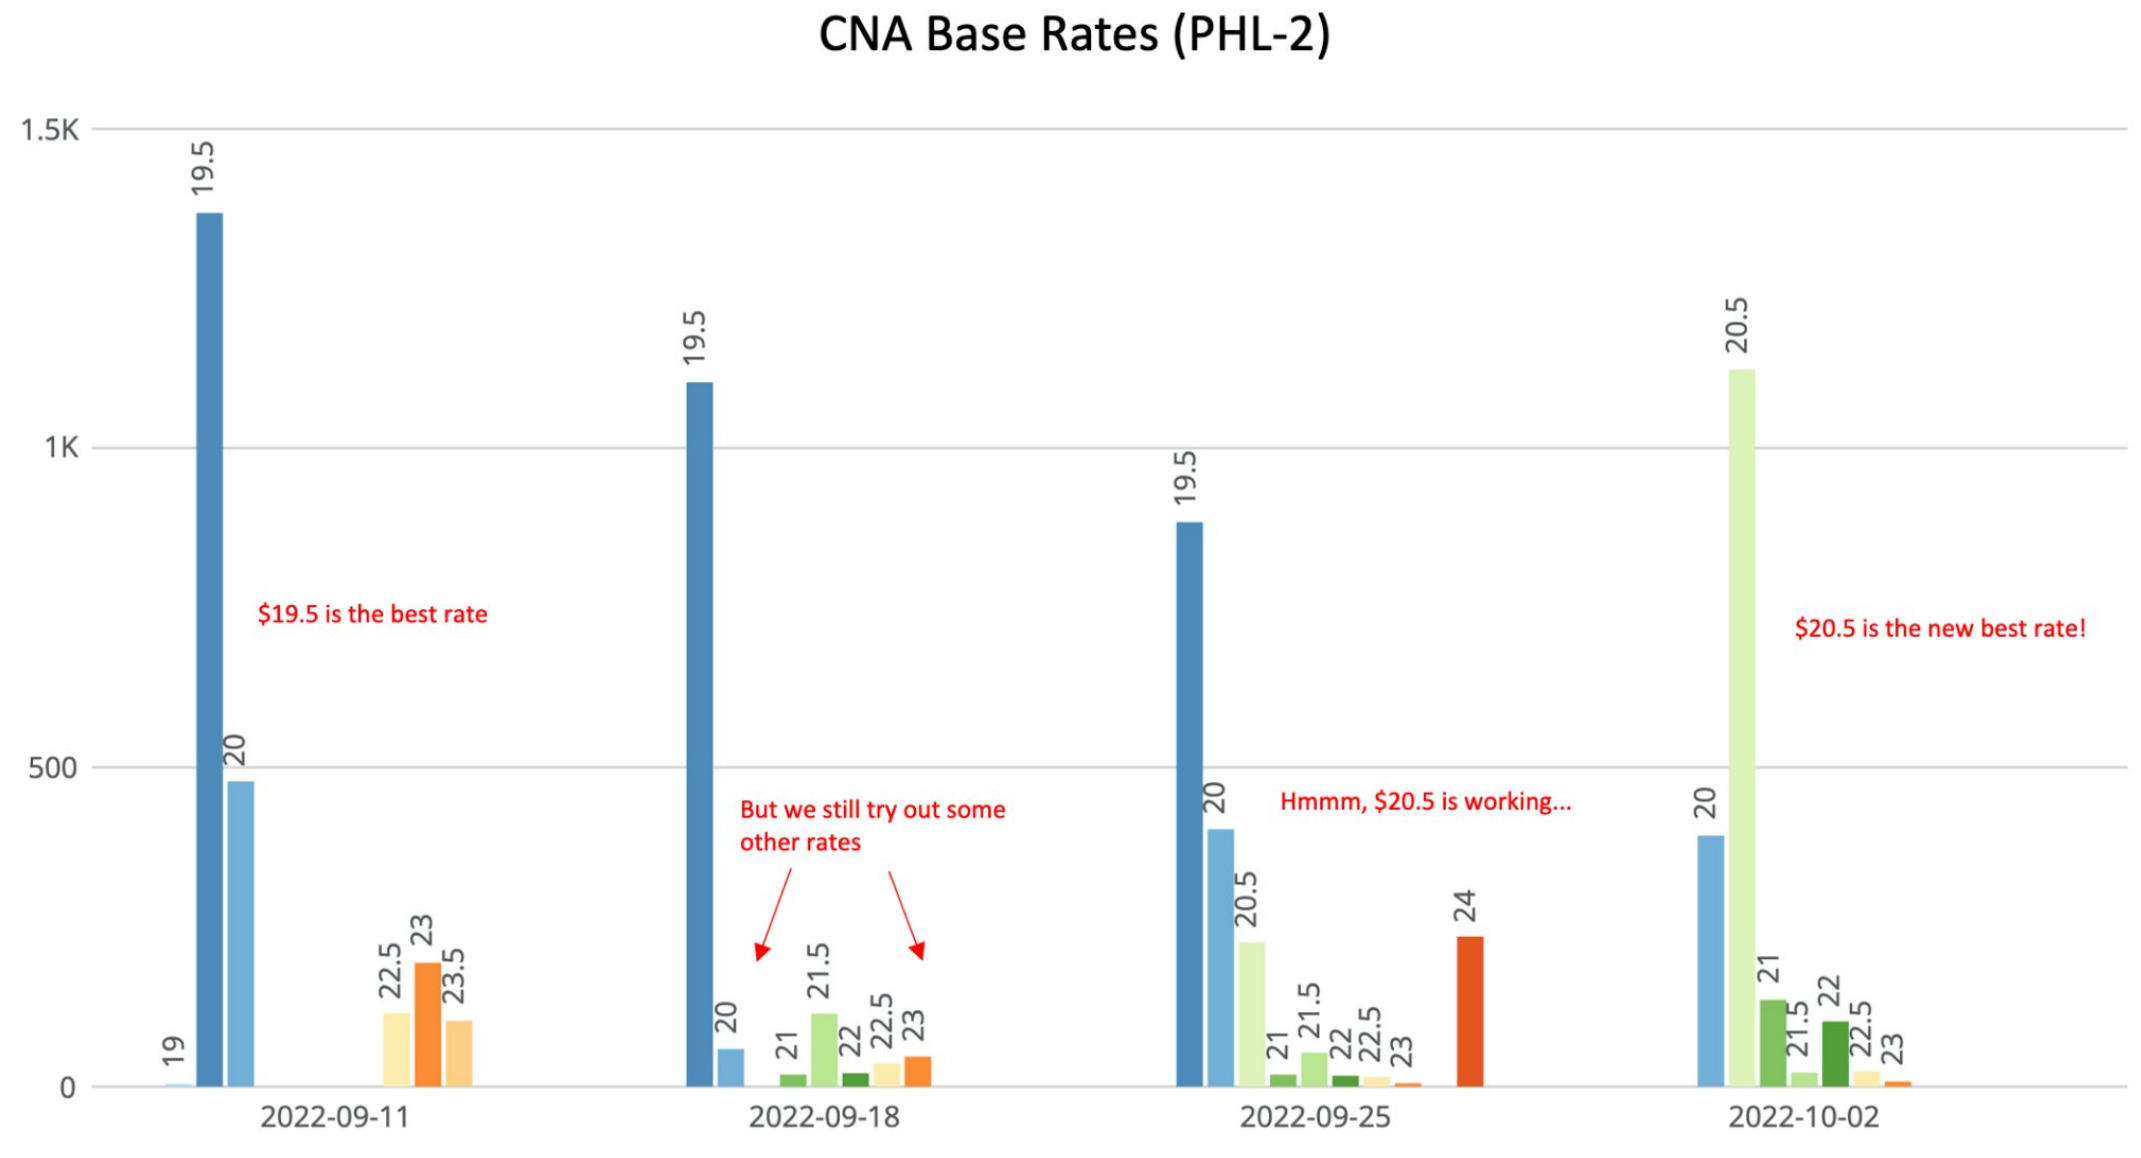

How did it go?

- Early 2019: We gradually roll out bandit pricing (previously a simple elasticity-based model)
- Late 2019: Product team releases "boost button" – facility can choose to pay $X/hr more for a given shift (pricing is now a 2-player game)
- Spring 2020: Covid pandemic → major increase in boost usage and overall demand
- Summer 2020: Big experiment - we put half our business on a industry-standard, fixed-margin pricing algorithm and keep the other half on our bandit model
- Spring 2022: Covid pandemic winds down → major decrease in boost usage and overall demand

How did it go?

- Summer 2020: Big Experiment - put half our business on a industry-standard, fixed-margin algorithm
    - The bandit model learned to anticipate facility actions → generated a degree of "boost sharing"
        - Facilities with low boost usage → bandit gives slightly higher initial rates
        - Facilities with high boost usage → bandit gives slightly lower initial rates
    - Allows us to fill both the easy/boosted shifts *and* the difficult unboosted shifts
    - Compared to the fixed-margin algorithm
        - ~10% higher profit dollars
        - Fewer manual corrections (~3% → less than 1%)

Try it at home?

- Maybe not ideal for long/thin catalogs
- Guardrails are important!
- For pricing, you might consider imposing monotonicity - our experience is the difference is small
- Robust to a lot of endogenous concurrent systems

Thank you

- 2019 Pricing Team: Michael Yang, Colin Shea
- Many other IntelyCare colleagues
- The CDSM Organizing Committee

Disclaimer - All opinions/errors are my own.

Email: btengelsen@intelycare.com
Social: https://www.linkedin.com/in/btengels/