Gurobi

Gurobean
THE COFFEE OPTIMIZATION GAME

Teaching Deck

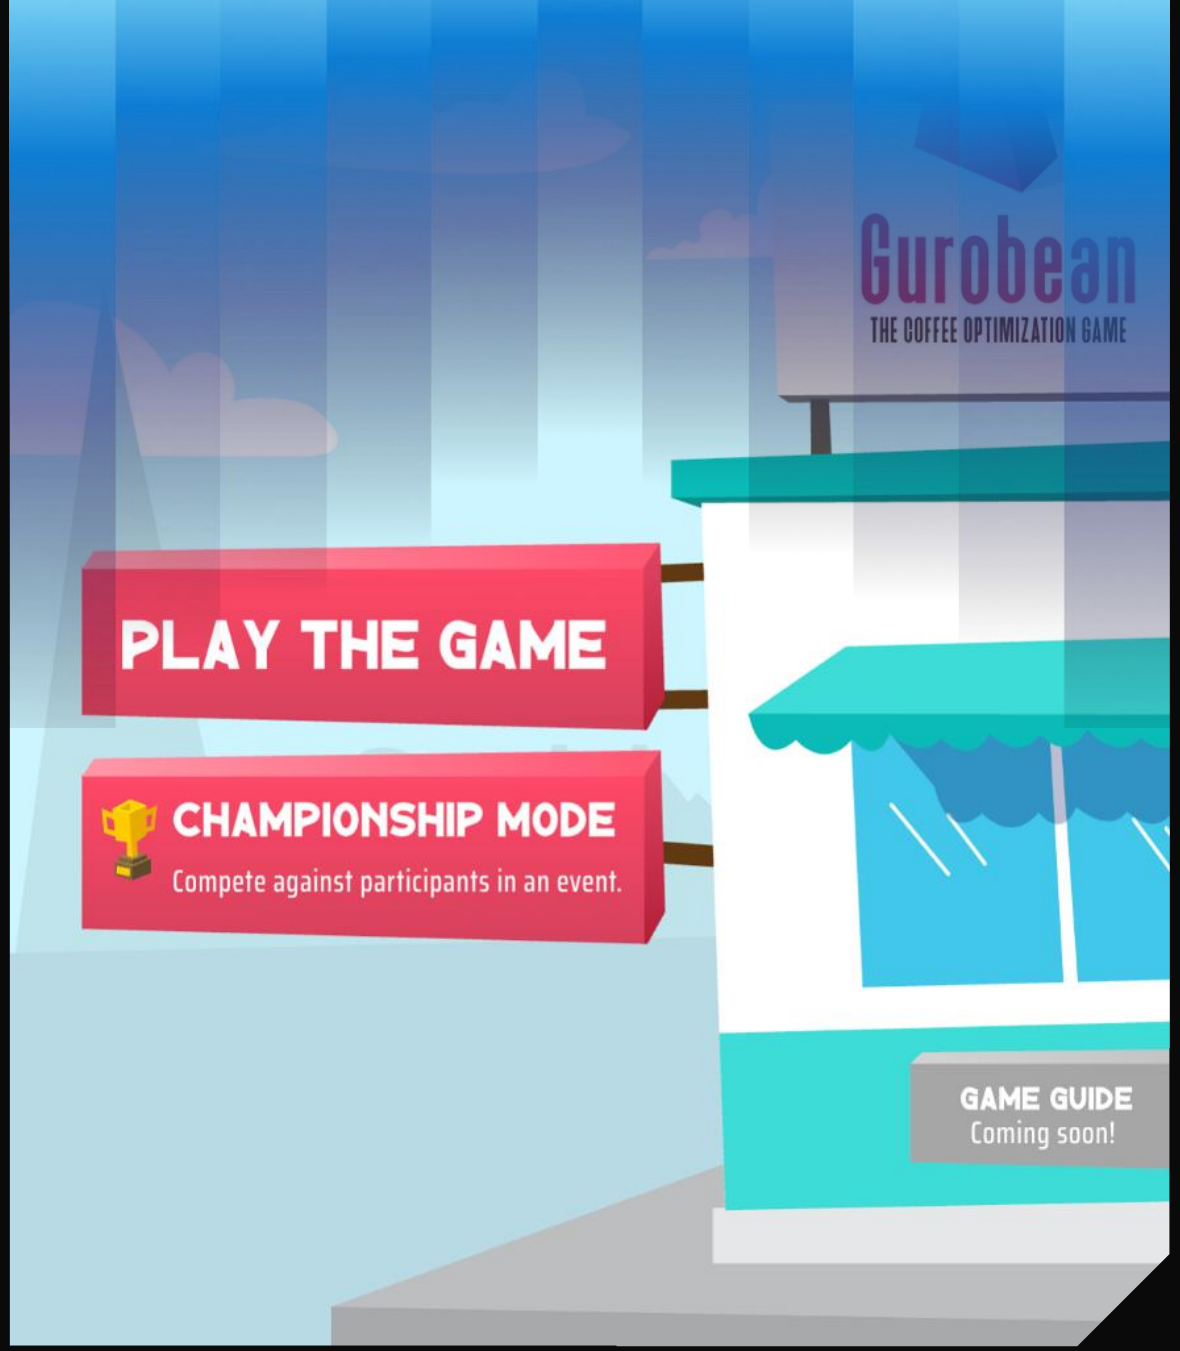

WWW.GUROBI.COM

Behind Every Great Decision

How to Use This Deck (for TAs)

One deck, two teaching paths. Every content slide is tagged in the top-right corner, hide the slides that don't fit your audience.

CORE: Everyone

TECHNICAL: OR / ENG / CS

BUSINESS: MBA / Analytics

Technical Track · 45–75 min

- Show: CORE slides + all TECHNICAL slides.
- Hide: BUSINESS slides.
- Emphasis: formulation, classification, why enumeration fails, deterministic vs. stochastic vs. robust.
- Good lead-in to integer/nonlinear programming, simulation, queueing.

Business & Analytics Track · 45–60 min

- Show: CORE slides + all BUSINESS slides.
- Hide: TECHNICAL math slides.
- Emphasis: decision quality, trade-offs, why feasible plans differ, the moving sweet spot.
- No math background needed, focus on insight and strategy.

Game Essentials

Core

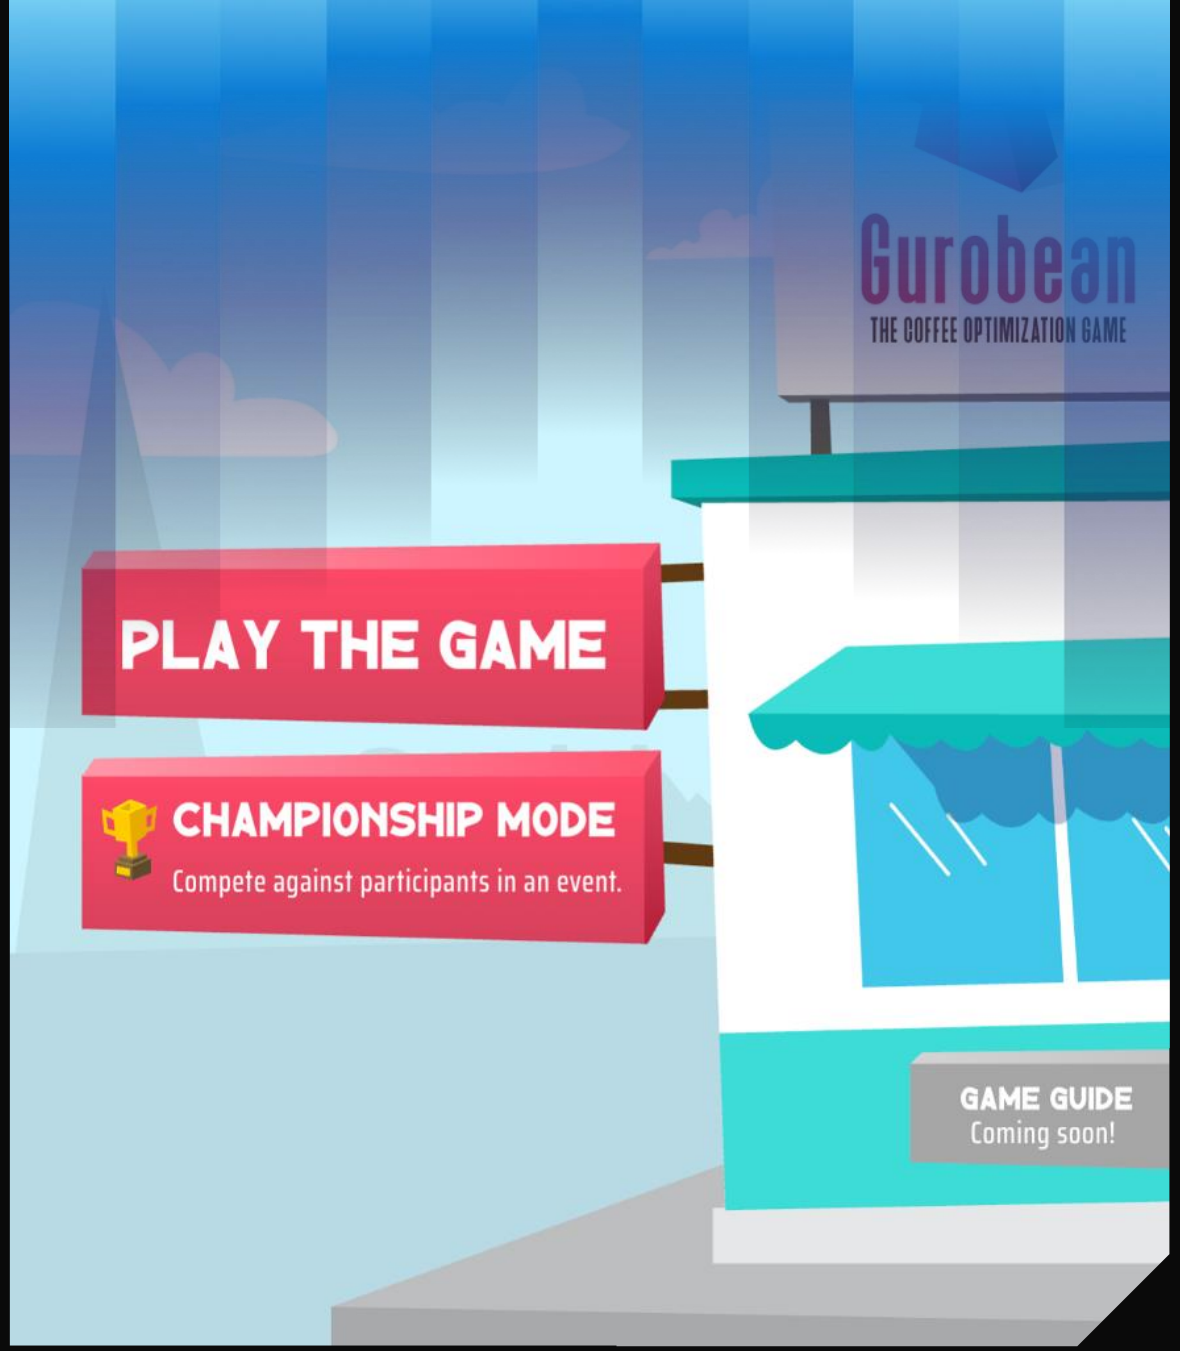

Gurobi

Gurobean needs your help!

Gurobi just opened a coffee shop

Business is already picking up!

- How much coffee should you brew each hour?
- How should you balance hot and cold drinks?
- What price should you charge, and how fast should the barista serve?

Your goal: run the shop to maximize 5-day profit.

Gurobi

CORE

The Coffee Shop

You run Gurobean in real time and watch the shop respond to your decisions as the day unfolds.

- **On the shop floor:** customers arrive and queue for hot or cold coffee, the barista serves one at a time, and BREWROBI brews to the rates you set.
- **Around it:** the top bar, live dashboard, and monitors track what's happening and how your plan is performing.

The next few slides tour each part of the screen.

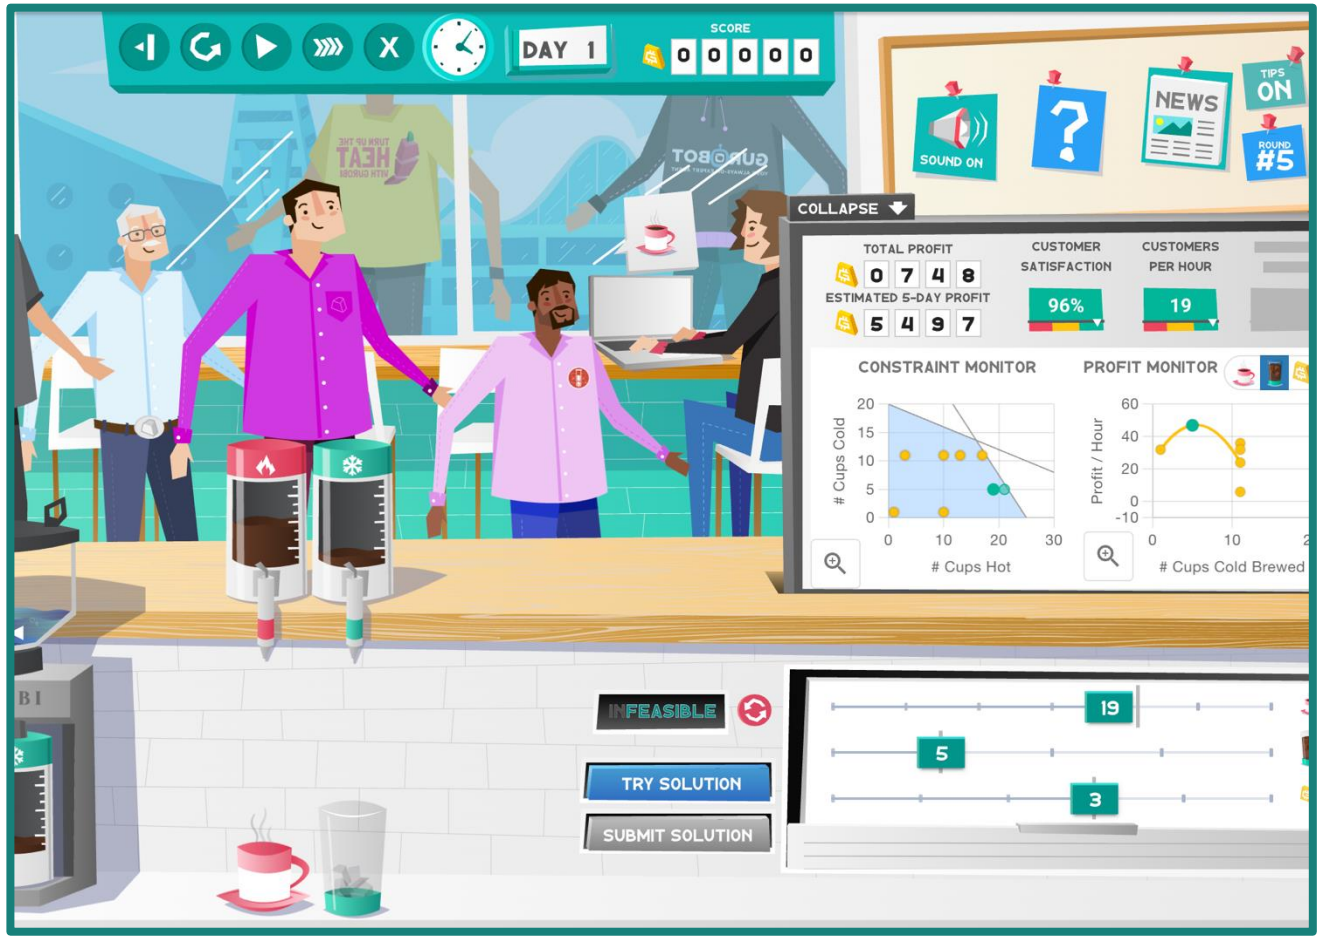

Gurobi

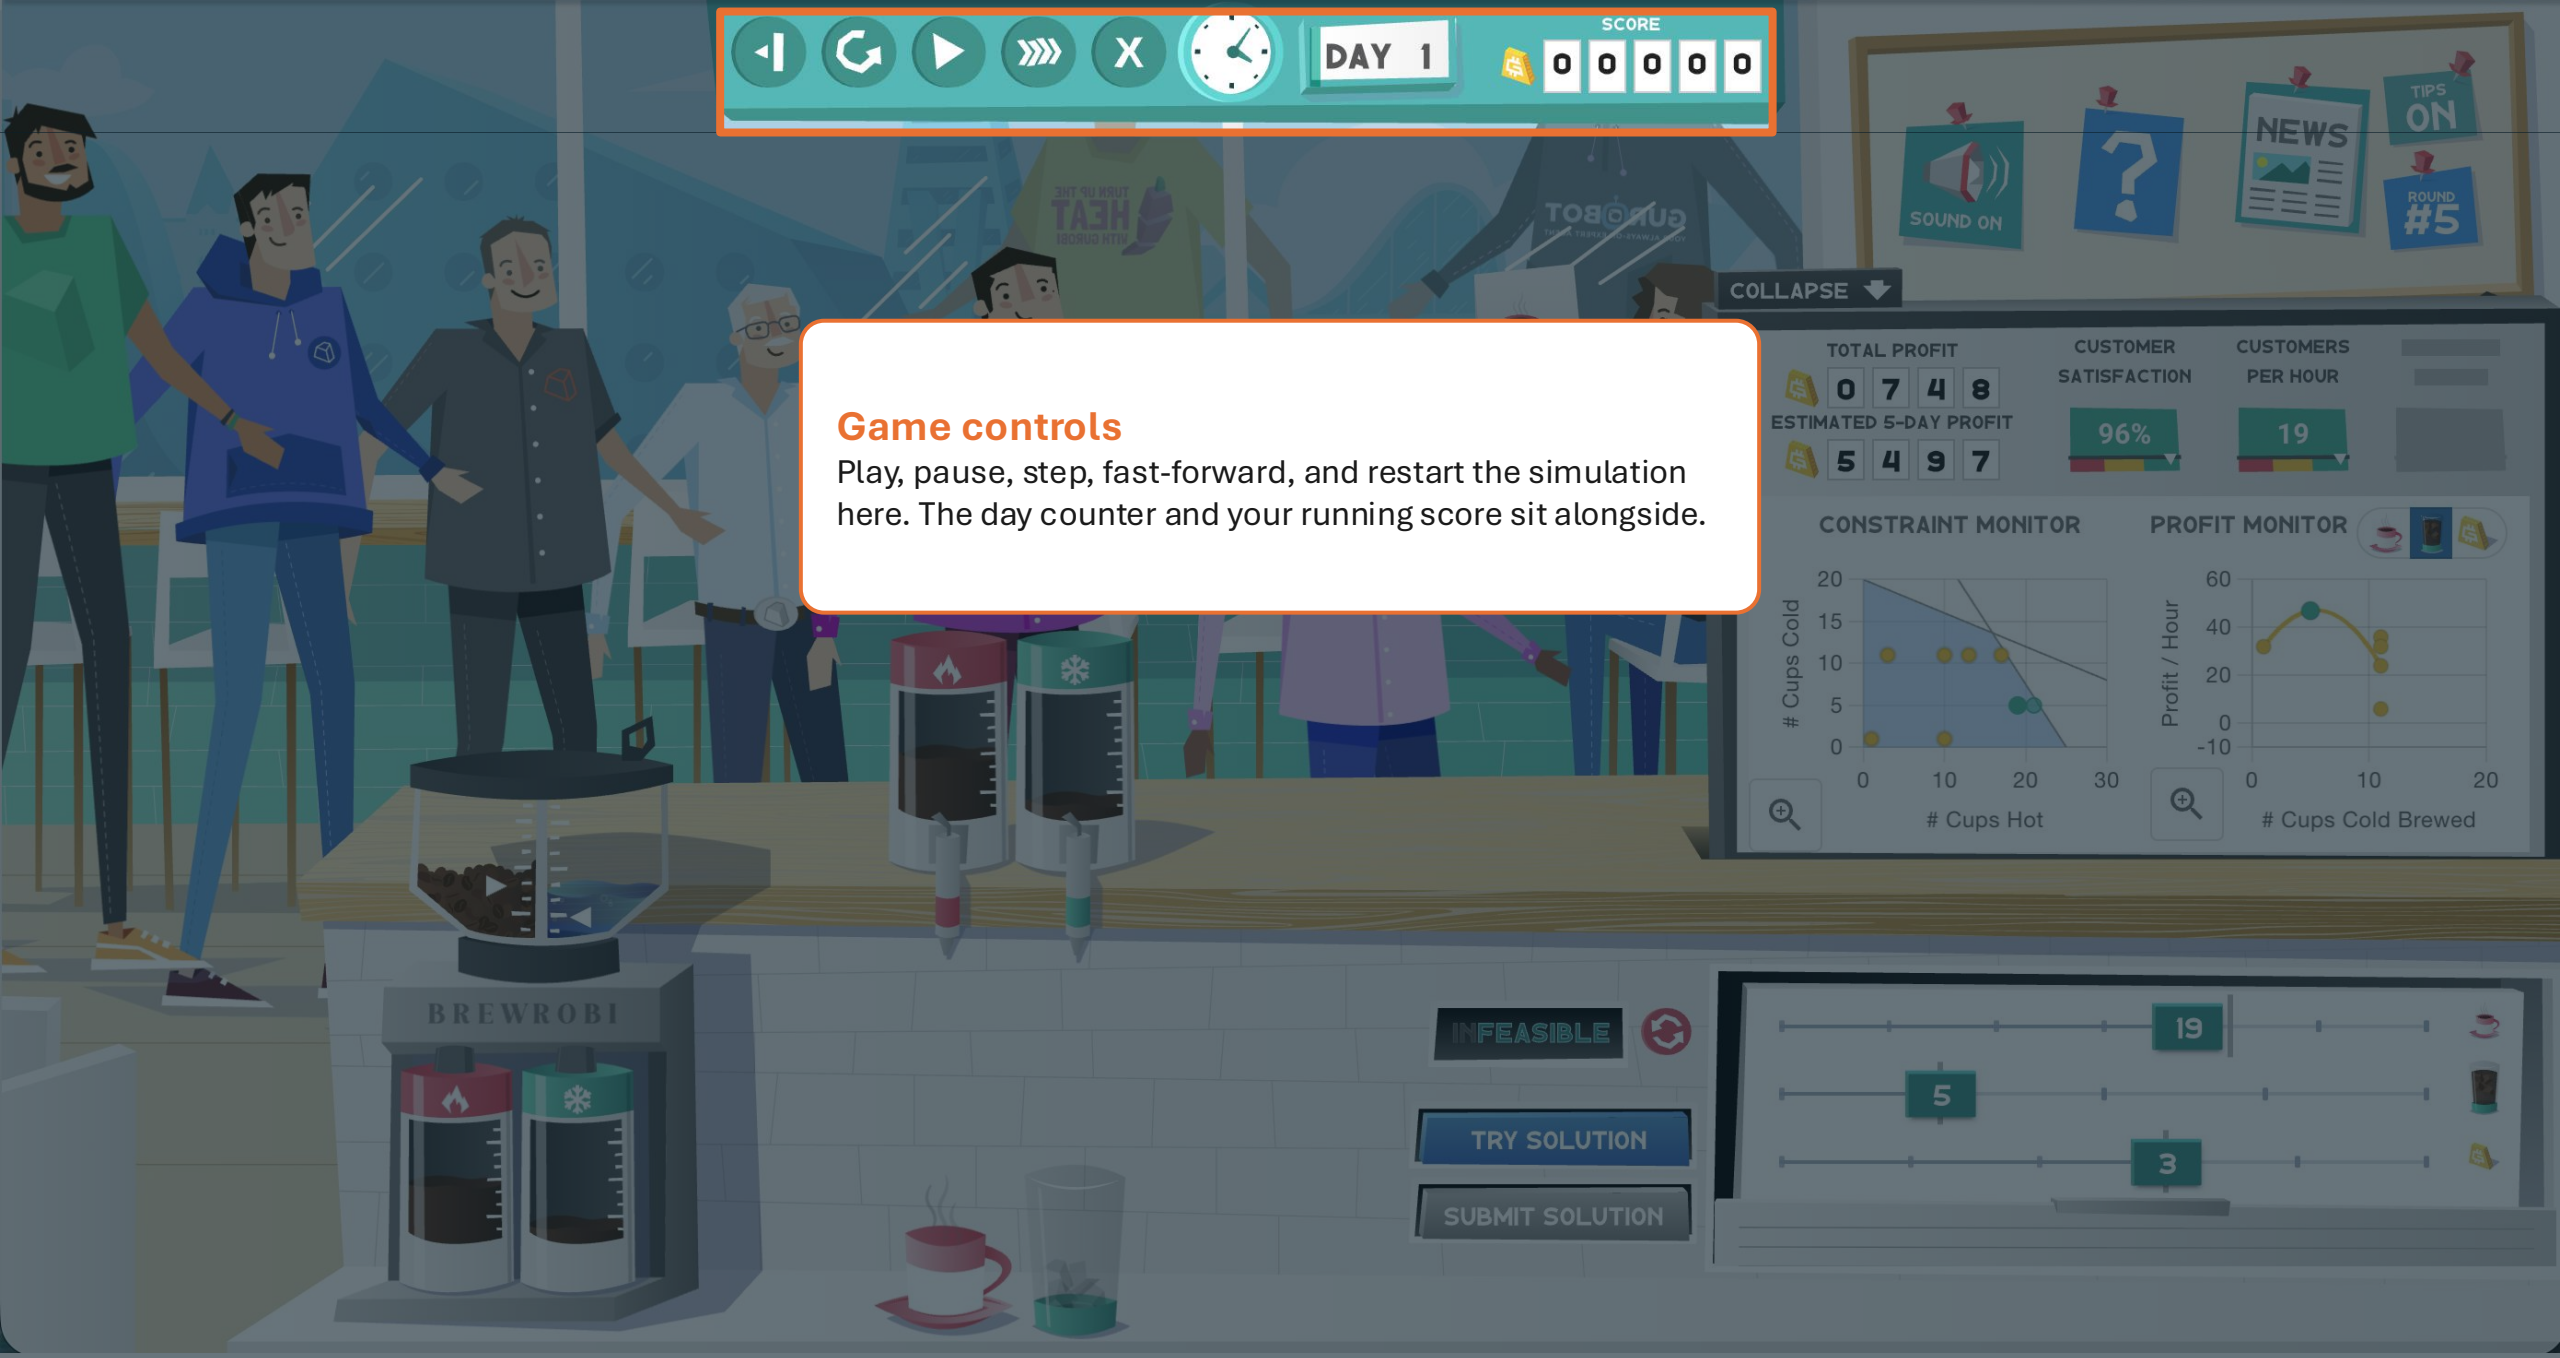

Game controls

Play, pause, step, fast-forward, and restart the simulation here. The day counter and your running score sit alongside.

The shop in action

Customers arrive at random, line up, and a single barista serves them one at a time. Watch the queue grow when demand outpaces service.

BREWROBI

The brewing station. Your brew-rate sliders decide how much hot and cold coffee is ready.

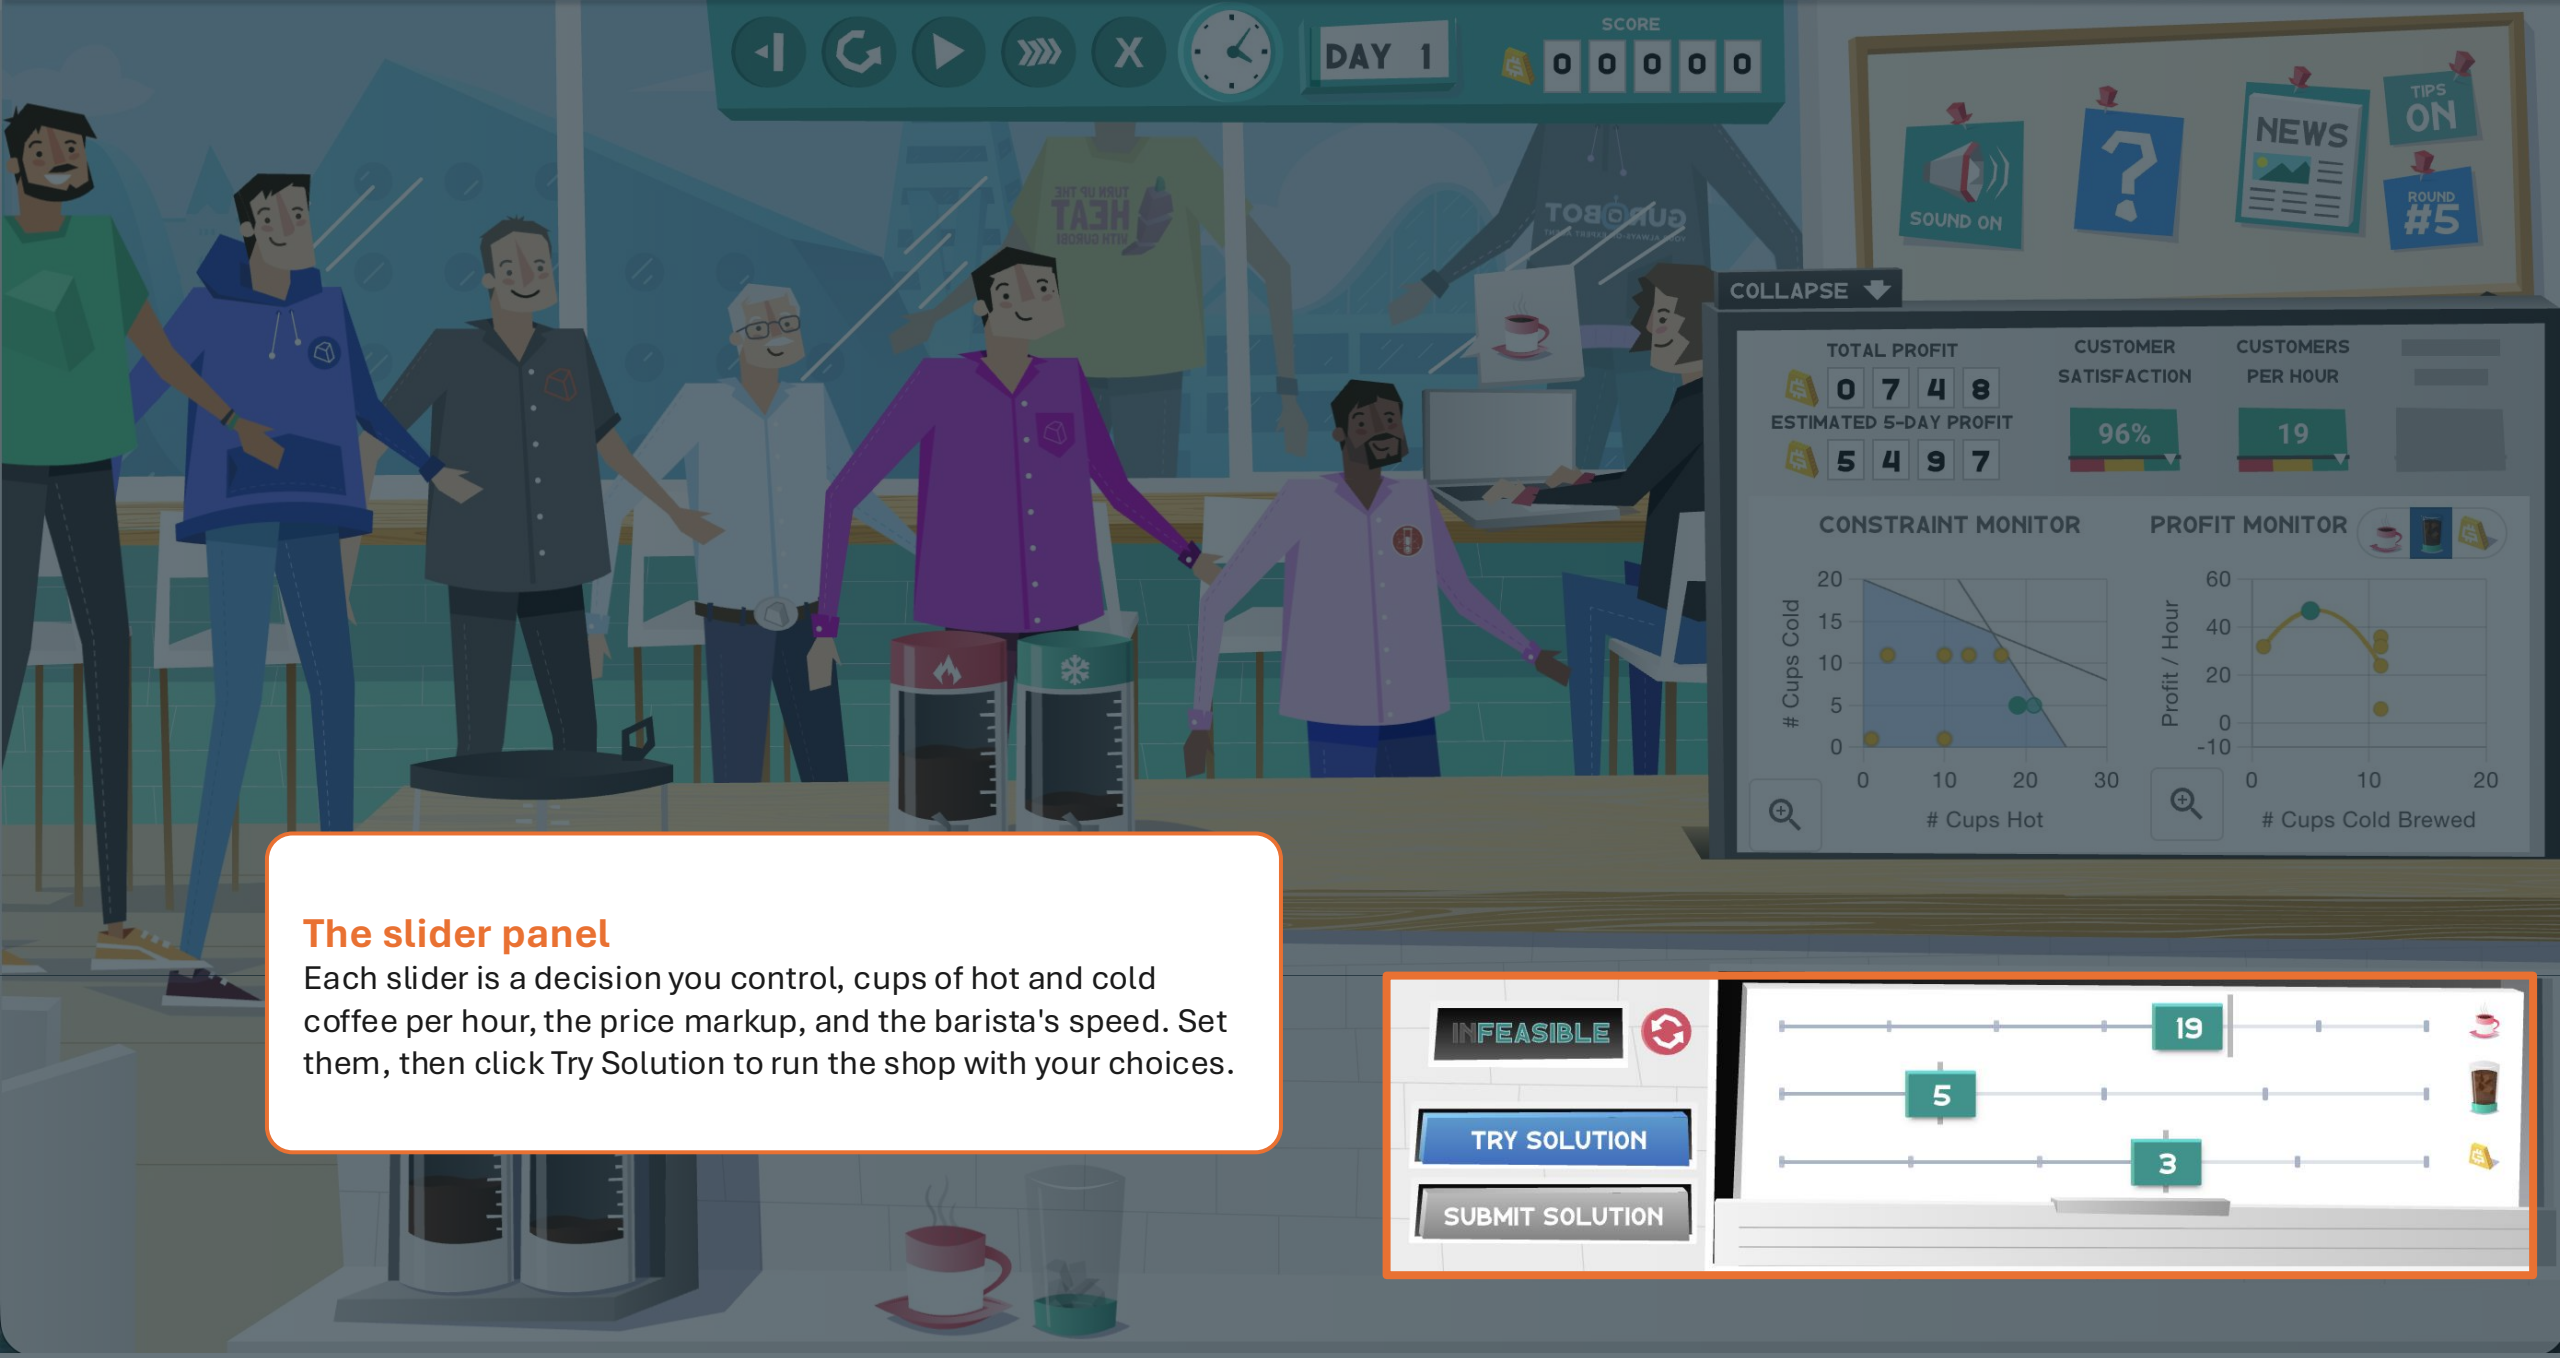

The slider panel

Each slider is a decision you control, cups of hot and cold coffee per hour, the price markup, and the barista's speed. Set them, then click Try Solution to run the shop with your choices.

CORE

Making Your Decisions

- Each slider is a decision variable you control, here, cups of hot and cold coffee to brew per hour, plus the price markup. Later rounds add barista service rate.
- Adjust the sliders to set a candidate solution, nothing changes until you act on it.
- Click **TRY SOLUTION** to restart the simulation with your new choices and watch it play out.
- If a setting needs more beans or water than the shop has, the panel reads **INFEASIBLE** and the values turn red, adjust until it's feasible (green).
- When you're happy, click **SUBMIT SOLUTION** to lock in your decisions and score the round.

Feasible, values are green

Infeasible, values turn red

Gurobi

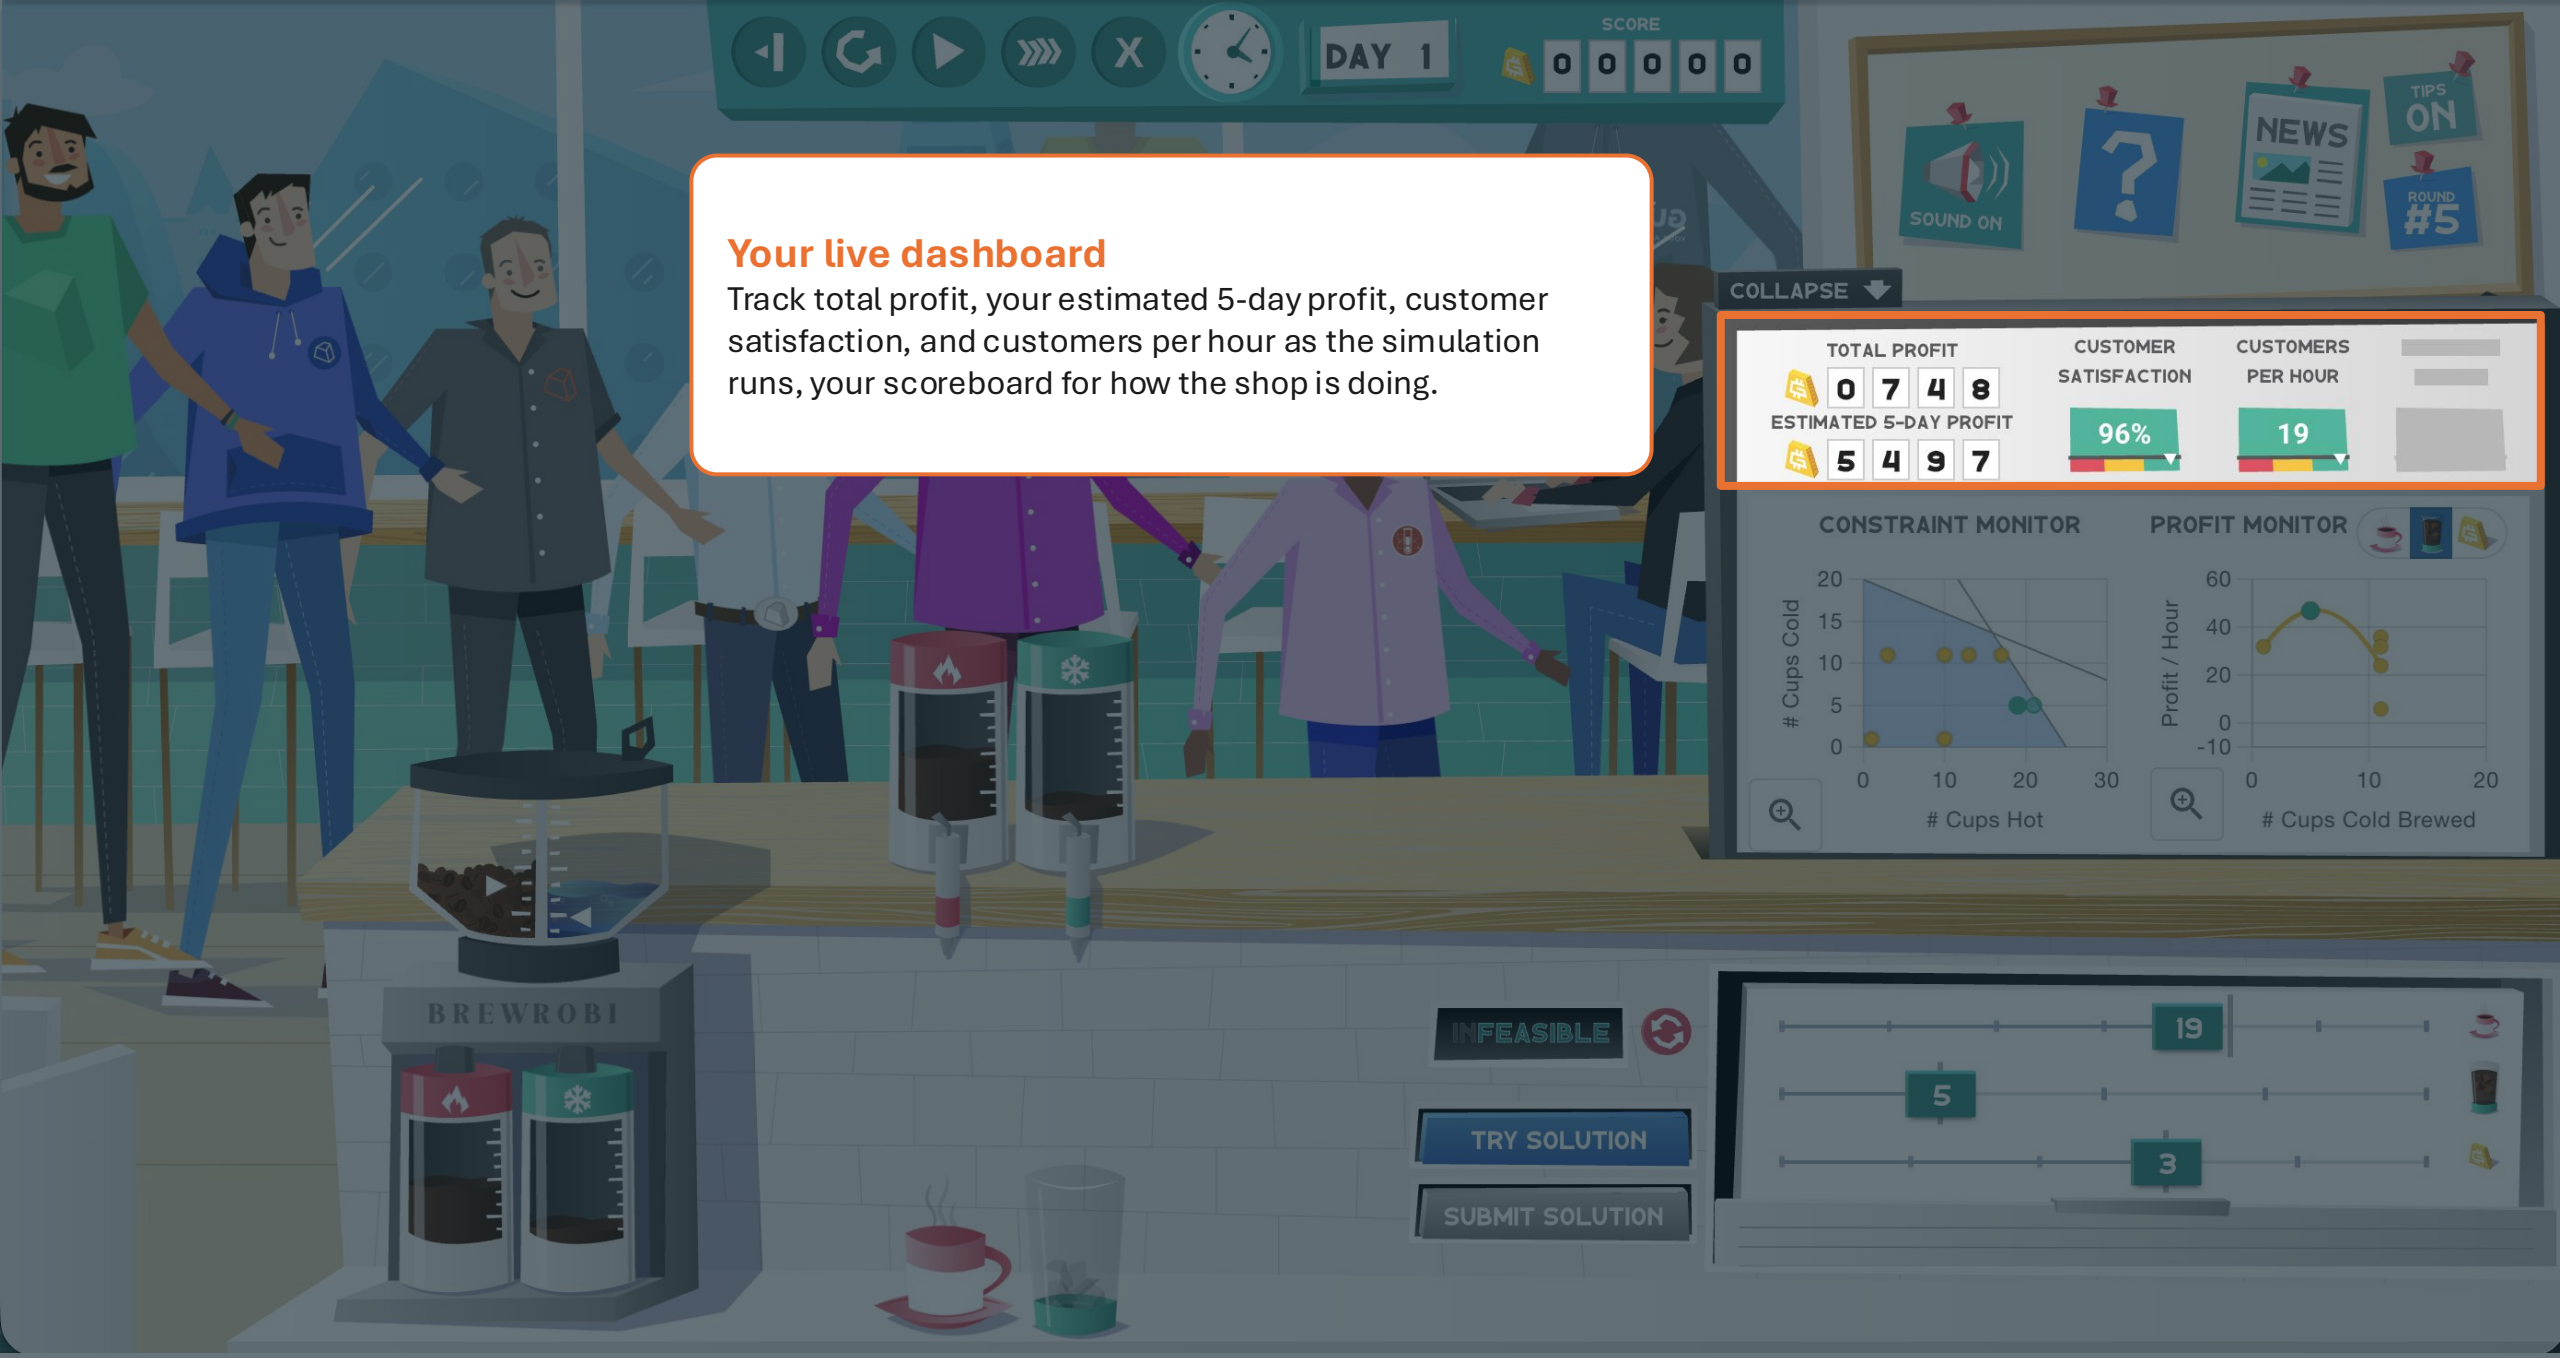

Your live dashboard

Track total profit, your estimated 5-day profit, customer satisfaction, and customers per hour as the simulation runs, your scoreboard for how the shop is doing.

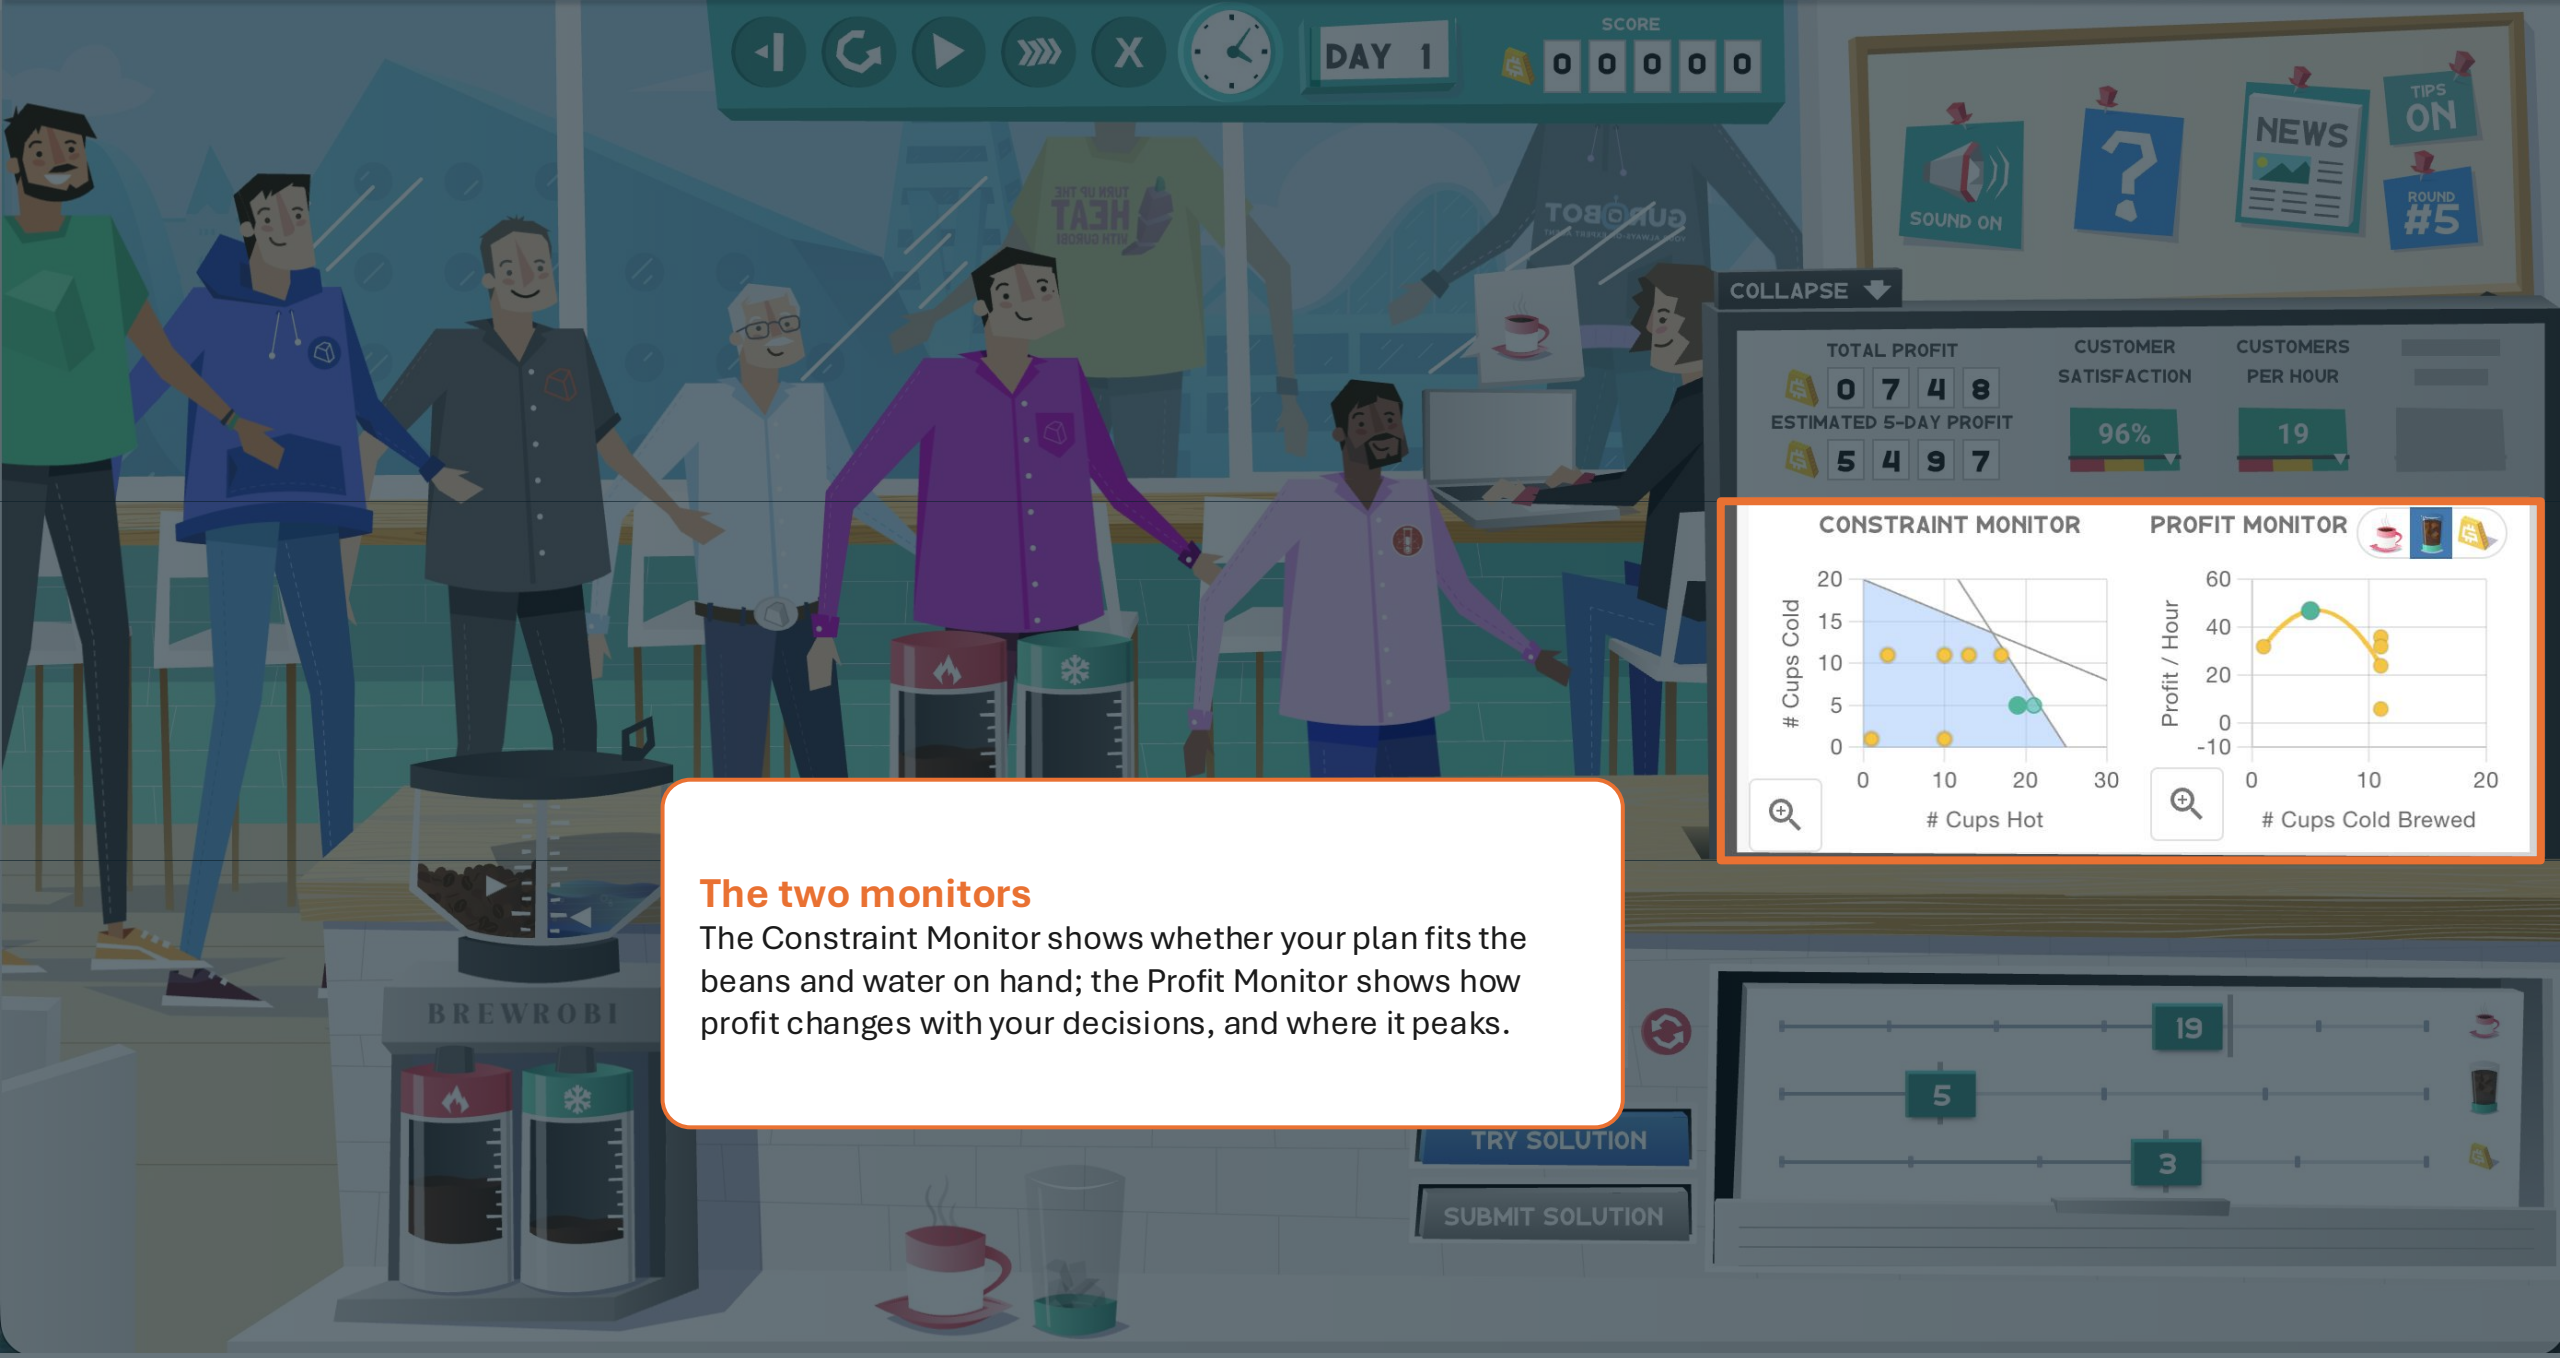

The two monitors

The Constraint Monitor shows whether your plan fits the beans and water on hand; the Profit Monitor shows how profit changes with your decisions, and where it peaks.

CORE

Reading the Monitors

Two dashboards turn the running simulation into decisions you can reason about.

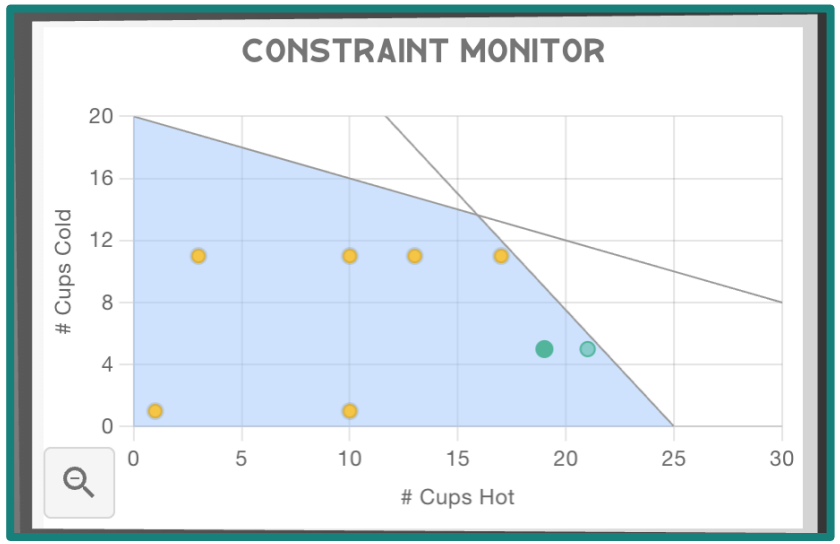

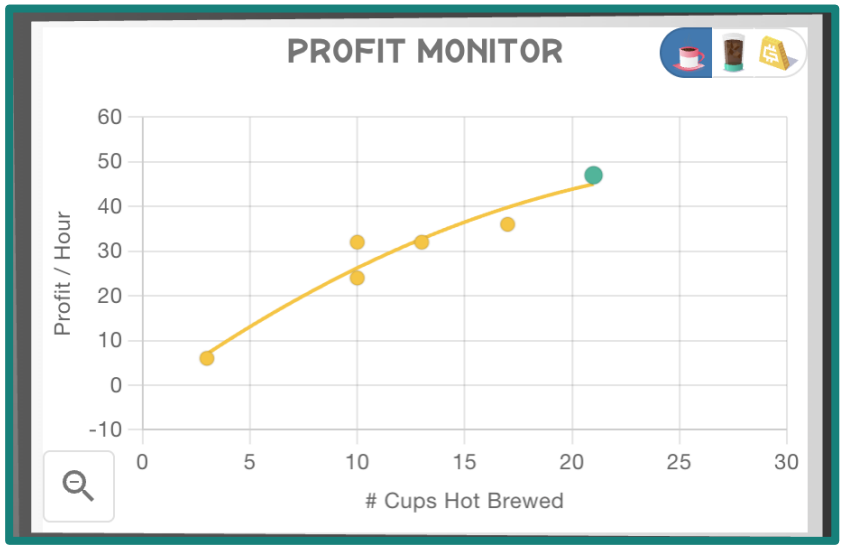

Constraint Monitor

- Shows whether a decision fits the shop's beans & water.
- Shaded area = the feasible region; markers outside it turn the sliders red.
- Answers: is this decision allowed?

Profit Monitor

- Plots profit against a decision (hot, cold, or price).
- The curve rises to a peak, then falls, brewing more isn't always better.
- Answers: is this decision good?

The bulletin board

Sound, help, the newsfeed, tips, and the round marker live here. At the start of each round, the newsfeed pops up to explain what's new.

Tips for Testing

CORE

Treat each round as an experiment: try several scenarios before you submit, and give each one time to play out.

- **Try, run, repeat**. Move the sliders, click Try Solution, and let each run play for at least 12 hours of sim time so its dot appears on the Profit Monitor.
- **Give it time to settle.** It starts empty and early numbers swing wildly; only once the shop reaches steady state do the metrics reflect your decision.
- **Enlarge the Profit Monitor.** As rounds get more complex, click the magnifier to expand it and read the curve clearly (shown right).
- **Cycle the x-axis**. Switch the Profit Monitor's x-axis between decision variables (hot, cold, price) to see how profit responds to each (shown right).

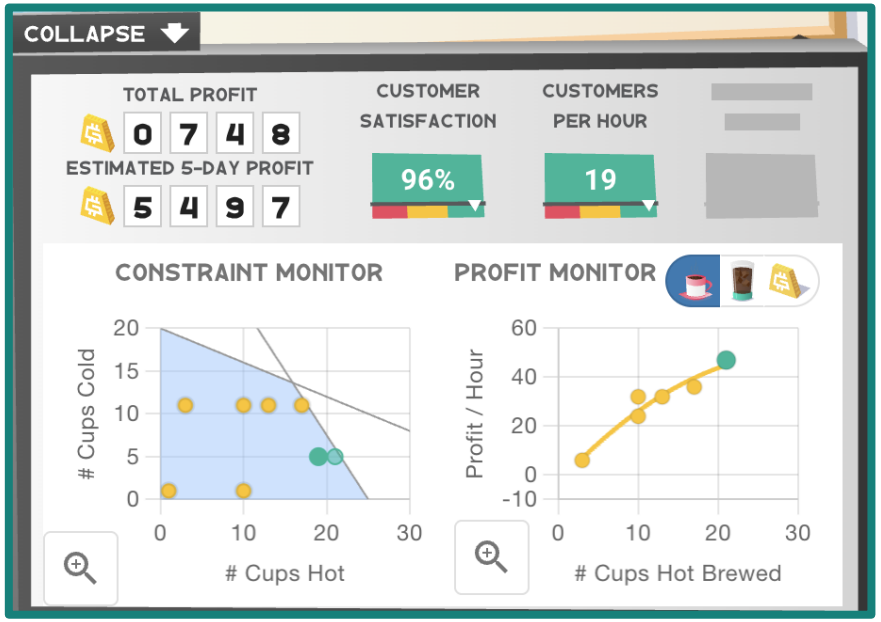

Gurobi

CORE

How the Rounds Work

Eight rounds, each adding one new wrinkle

The game has eight rounds of rising complexity. Each one starts with a newsfeed pop-up (right) where Alinea introduces the new twist and shows that round's parameters.

- **Rounds 1–2:** brew hot, then hot + cold
- **Rounds 3–4:** brewing now costs money
- **Round 5:** set the price (markup)
- **Round 6:** crowds balk and leave
- **Round 7:** customers order multiple cups
- **Round 8:** choose the barista's speed

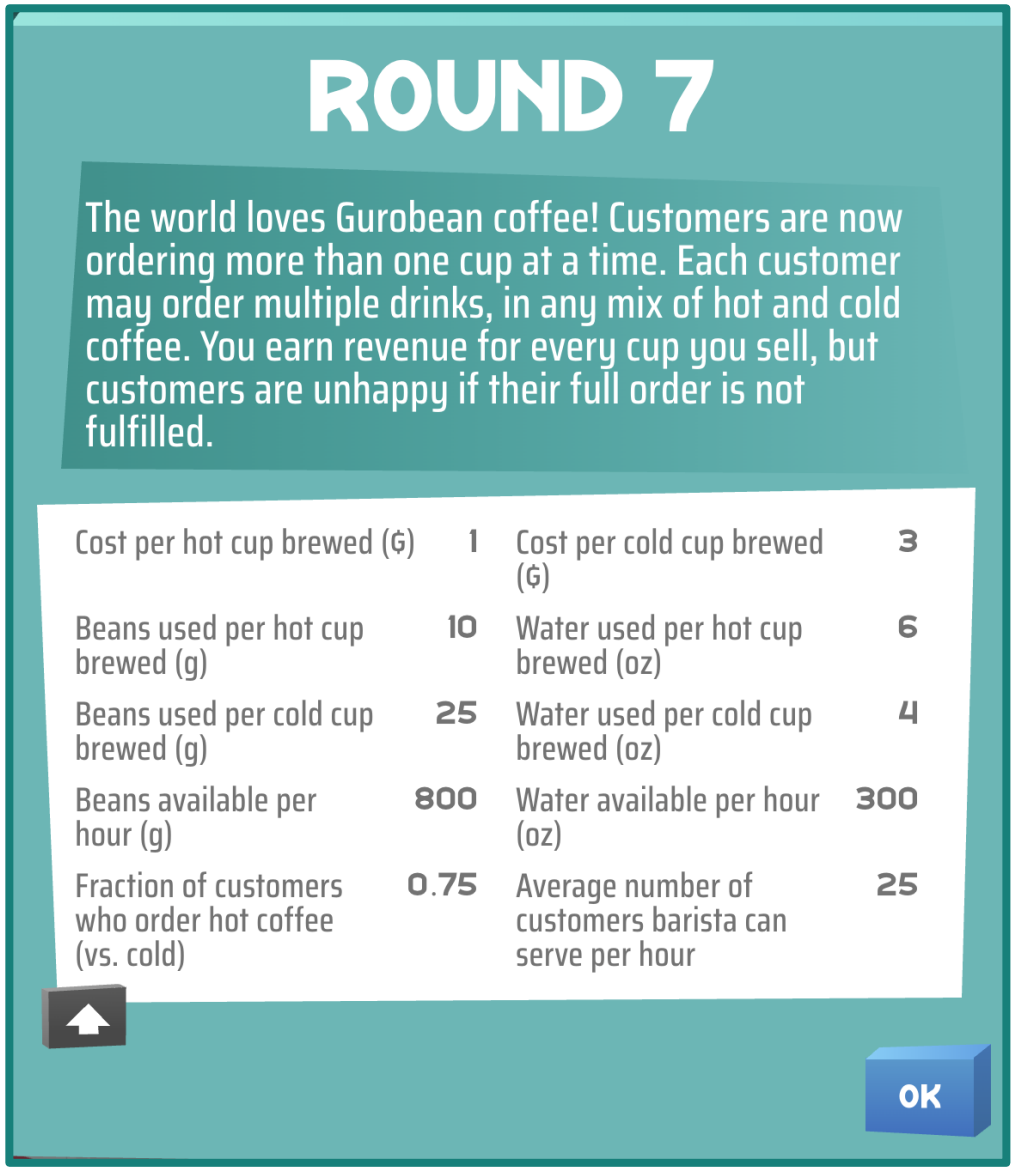

CORE

You Decide Before You Know Demand

When you decide

- You set brew rates, price, and service before the day starts.
- All you have is the average arrival rate, the forecast.
- Brew too little and you miss sales; too much and cups go to waste.

What you're graded on

- Customers actually arrive at random (a Poisson process).
- Your 5-day profit is scored on what really happened, not the forecast.
- So the same decision yields a different result every run.

This gap between planning on forecasts and being judged on reality is what makes a good decision hard to recognize.

Demand splits two ways, too: the newsfeed gives the fraction of customers who want hot vs. cold (e.g. 0.75 hot), so you must forecast, and commit to brewing, each type separately.

Gurobi

CORE

You vs. Gurobi

When you Submit, the game scores your decision with a full five-day simulation, then Gurobi solves the round and the two are compared side by side:

- You see your slider settings and resulting metrics next to Gurobi's optimal settings.
- Each solution is plotted on the expected-profit curve, filled markers show expected profit, hollow markers show what actually happened in this run.
- Read the two markers together: because this run is random, your hollow marker can land above or below Gurobi's, the outcome alone doesn't tell you whose decision was better.

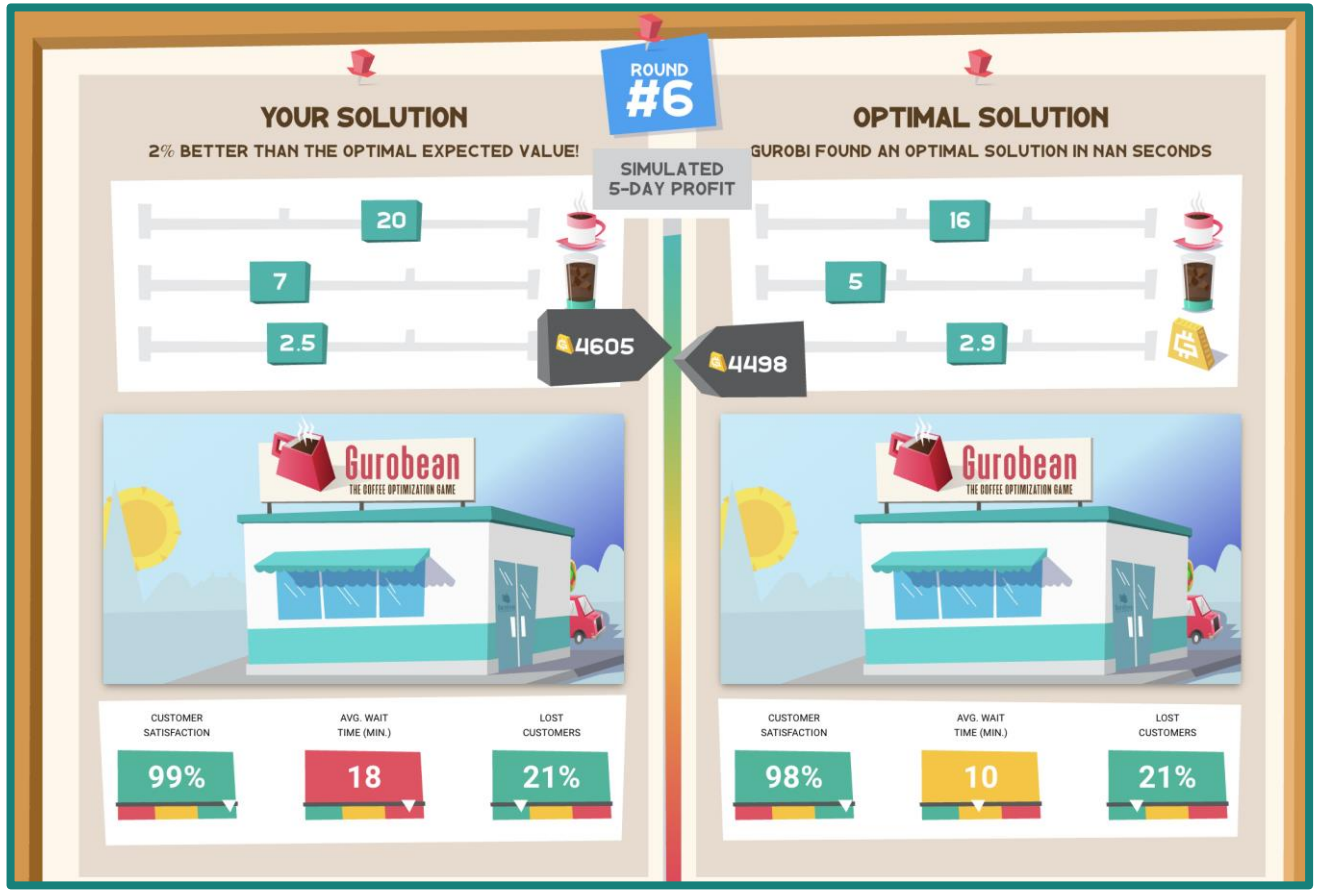

Gurobi

CORE

Expected vs. Realized

The central lesson of the comparison screen

- A single five-day simulation is just one realization of a random system.
- Gurobi optimizes expected profit, the decision that performs best on average across many possible futures.
- So you can beat Gurobi on one run and Gurobi's decision is still the better one.
- **The real question isn't “who got lucky this time?” but “which decision would you trust before knowing what randomness will occur?”**

Gurobi

CORE

Activity: Predict, Then Run

Before you touch Try Solution, make a prediction.

1. Pick one slider, say, raise the price (markup).
2. Ask the room: what happens to profit? To the queue? To lost customers? Have students commit to an answer.
3. Now run it. Did the system behave the way you predicted, and did it play out the same way twice?

Why it works
Committing to a prediction first makes the surprise stick, and sets up the lesson that nonlinear, random systems don't move the way intuition expects.

Championship Mode

Two ways to compete on decision quality, not luck

Two ways to compete

- **Class Championship**, private, via a match code you share
- **Global Championship**, public, a fresh scenario each week

Why it's fair

- Everyone plays the same scenario, so nobody wins on a lucky day
- Rewards consistent, well-reasoned decisions

Technical Deep-Dive

Technical Track · OR / Engineering / CS

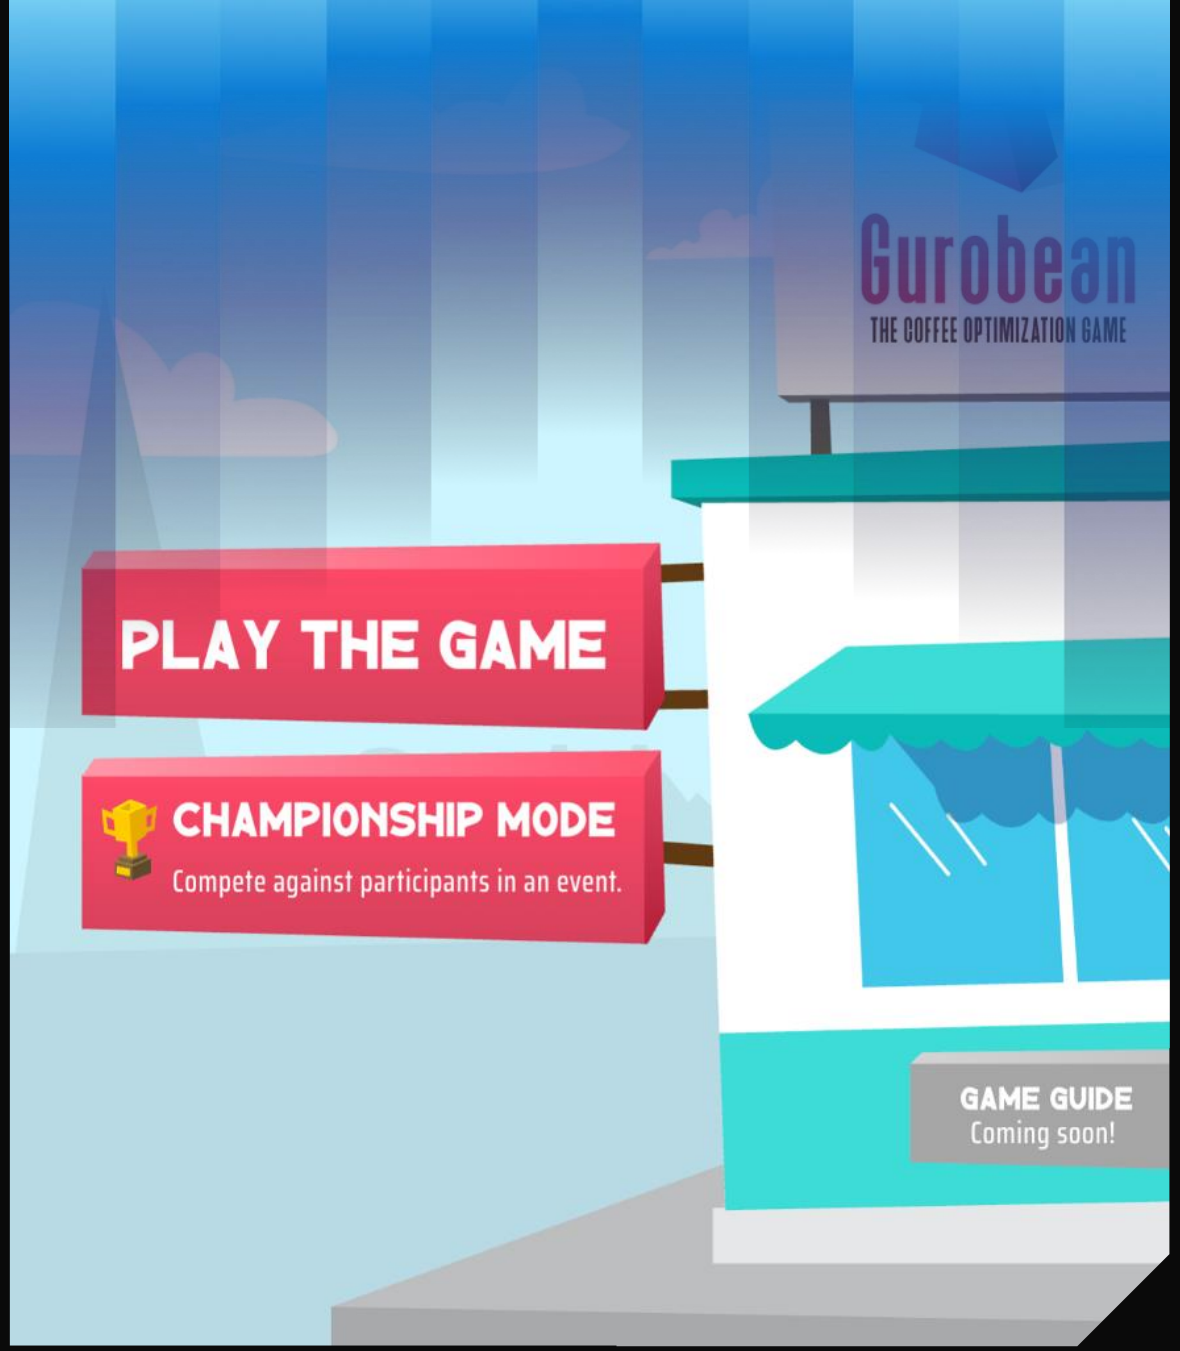

Gurobi

TECHNICAL

What Students Should Take Away

Technical learning outcomes, the lasting points behind the modeling and simulation

- **Models define decisions, objectives, and constraints, but evaluation may need simulation.** A formulation says what is optimal in theory; simulation shows how it behaves in practice.
- **An expected-value-optimal solution will not win every run.** The decision that is best on average can still lose on any single realization.
- **Nonlinearity and congestion create feedback that breaks simple reasoning.** Small changes ripple through queues and prices in ways intuition misses.
- **Simulation output is statistical, not deterministic.** Read outcomes as distributions, and let the system reach steady state before judging.
- **Feasibility defines the solution space; performance defines its quality.** Meeting every constraint is the entry ticket, not the goal.

TECHNICAL

Under the Hood

The Models Behind the Game

- Newsvendor-style inventory decisions
- **Stochastic demand** (random arrivals)
- **Resource constraints** (beans, water, capacity)
- Queueing theory
- **Nonlinear** profit functions
- **Simulation-based evaluation**

Gurobi

TECHNICAL

Activity: Formulate Round 1

Round 1: brew hot coffee only, limited beans and water

Write the three pieces:

- **Objective**, what are we maximizing?
- **Constraints**, what limits us?
- **Decision variable**, what do we choose?

Give students 2 minutes in pairs, then reveal →

One answer

maximize expected profit from hot coffee
subject to beans used ≤ beans available; water used ≤ water available
by choosing Q_h = cups of hot coffee to brew per hour

Same three pieces as every optimization model.

TECHNICAL

Notation

What each symbol in the model means

Decisions · what you and Gurobi choose

Q_h, Q_c: cups of hot and cold coffee brewed per hour

m: markup added to cost (a decision in later rounds)

μ: barista service rate (a decision in later rounds)

Demand

λ: customer arrival rate per hour

ψ_h, ψ_c: share of customers wanting hot / cold

D_h, D_c: random hourly demand (mean λ_i, Normal approx.)

Prices and costs

r_h, r_c: revenue per cup sold (hot / cold)

c_h, c_c: cost per cup brewed (hot / cold)

Resources

a_b, a_w: beans and water available

u_{hb}, u_{hw}, u_{cb}, u_{cw}: ingredient used per cup

Profit-curve terms

ϕ, Φ: standard normal density and CDF

$z_i = (Q_i - \lambda_i) / \sqrt{\lambda_i}$: standardized brew quantity

Gurobi

TECHNICAL

Nonlinear Model

Here's the base model Gurobi solves in the early rounds; you don't need to follow the math to play.

1 · Demand responds to markup

$$\lambda(m) = \bar{\lambda}\exp\left(\frac{am}{m_0}\right), \qquad a = \ln\left(\frac{\lambda_0}{\bar{\lambda}}\right)$$

$$\lambda_h = \lambda\psi_h, \qquad \lambda_c = \lambda\psi_c$$

2 · Expected profit per drink

$$P_i(Q_i) = (r_i - c_i)Q_i - r_i\,\mathbb{E}[(Q_i - D_i)^+]$$

$$P_i(Q_i) = (r_i - c_i)Q_i - r_i\sqrt{\lambda_i}\,\left[\phi(z_i) + z_i\Phi(z_i)\right], \qquad z_i = \frac{Q_i - \lambda_i}{\sqrt{\lambda_i}}$$

3 · The base model

$$\underset{Q_h, Q_c}{\text{maximize}} \quad P(Q_h, Q_c) = P_h(Q_h) + P_c(Q_c)$$

$$\begin{aligned} \underset{Q_h, Q_c}{\text{maximize}} \quad & P_h(Q_h) + P_c(Q_c) \\ \text{subject to} \quad & u_{hb}Q_h + u_{cb}Q_c \le a_b, \\ & u_{hw}Q_h + u_{cw}Q_c \le a_w, \\ & Q_h, Q_c \ge 0 \end{aligned}$$

Gurobi

TECHNICAL

The Shape of the Objective

Where the peak sits, how high, and how sharp

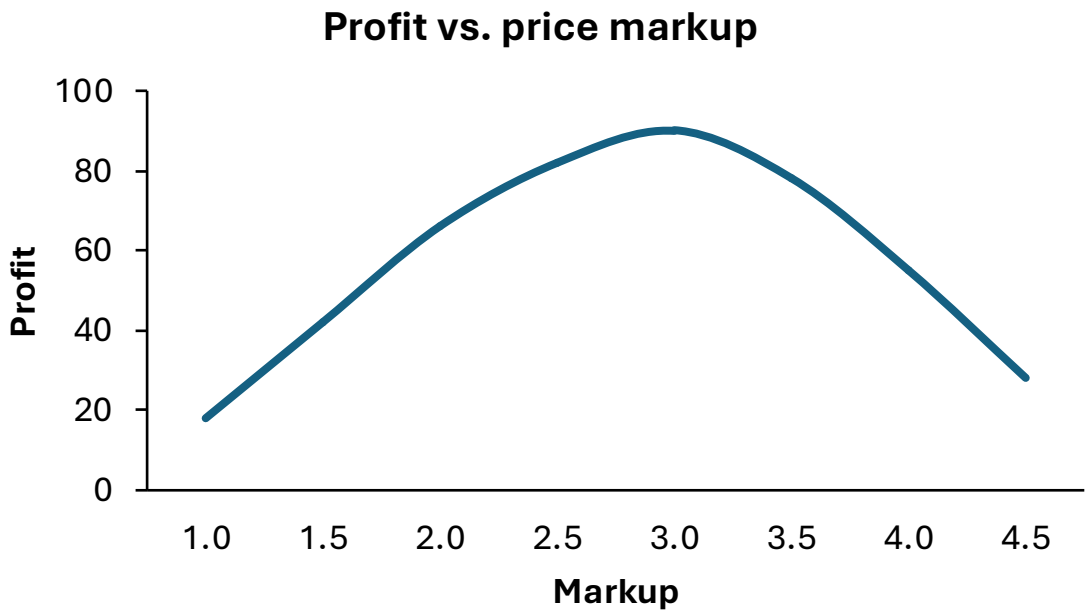

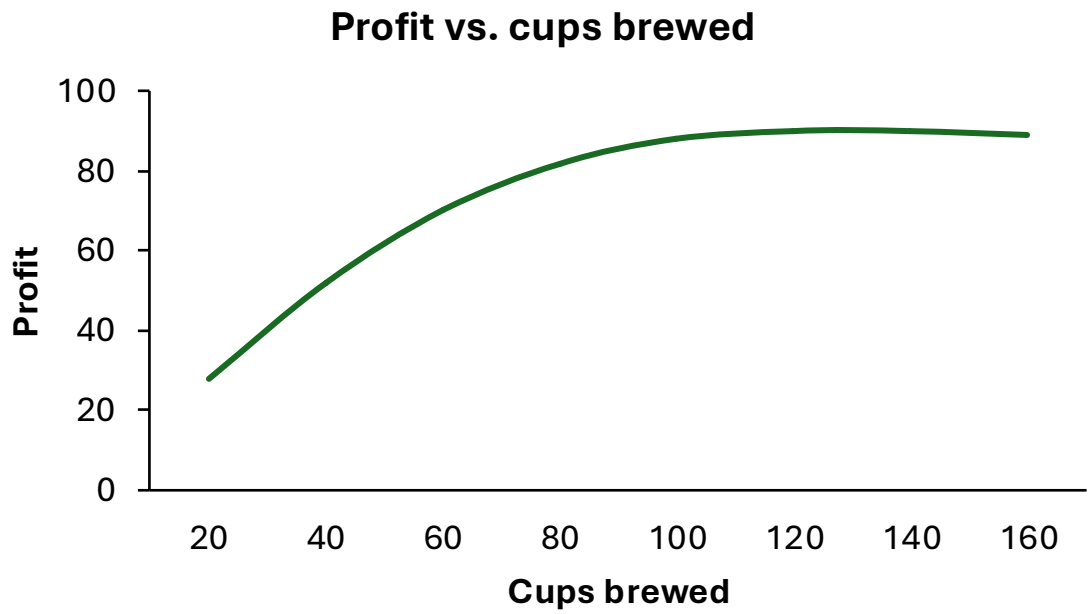

Flat peak: forgiving, so roughly right is good enough. **Sharp peak:** unforgiving, so a small miss is costly.

The best choice moves with each scenario, and a binding resource limit can pin it to the constraint, not the curve's natural peak.

TECHNICAL

Why We Simulate

The formula is only one piece; simulation scores the real shop

- **A formula covers only one piece.** The closed form approximates a single drink's profit; the real shop (queues, balking, congestion, randomness) is dynamic, so its true performance has no formula.
- **We estimate by running it.** Simulate the shop many times and average, the way call centers, hospitals, and traffic networks are analyzed.
- **The output is a distribution.** Read results statistically; any single run is just one draw, not the verdict.
- **Let it reach steady state.** Early numbers swing while the shop fills, so wait for burn-in before judging a decision.

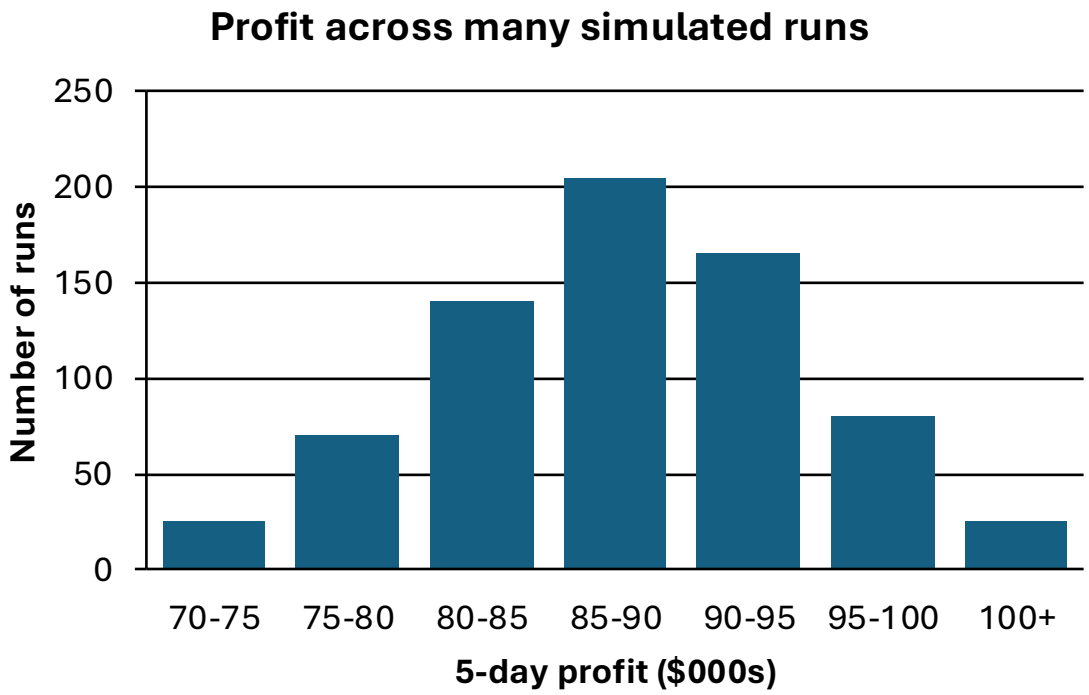

TECHNICAL

How Do You Actually Solve It?

- **By hand / trial-and-error**: What you did playing Gurobean. Fun for a while, but slow and rarely optimal.
- **Try every combination (enumeration)**: Conceptually simple, but the combinations explode into the trillions, hopeless.
- **Heuristics, smart rules of thumb**: Break the problem up, decide one slider at a time, or simplify it. Useful and fast, but approximate.
- **Optimization solvers**: Software like Gurobi uses math and algorithms to find the optimal decision in a fraction of a second.

Business & Strategy

Business & Analytics Track · MBA / Exec Ed

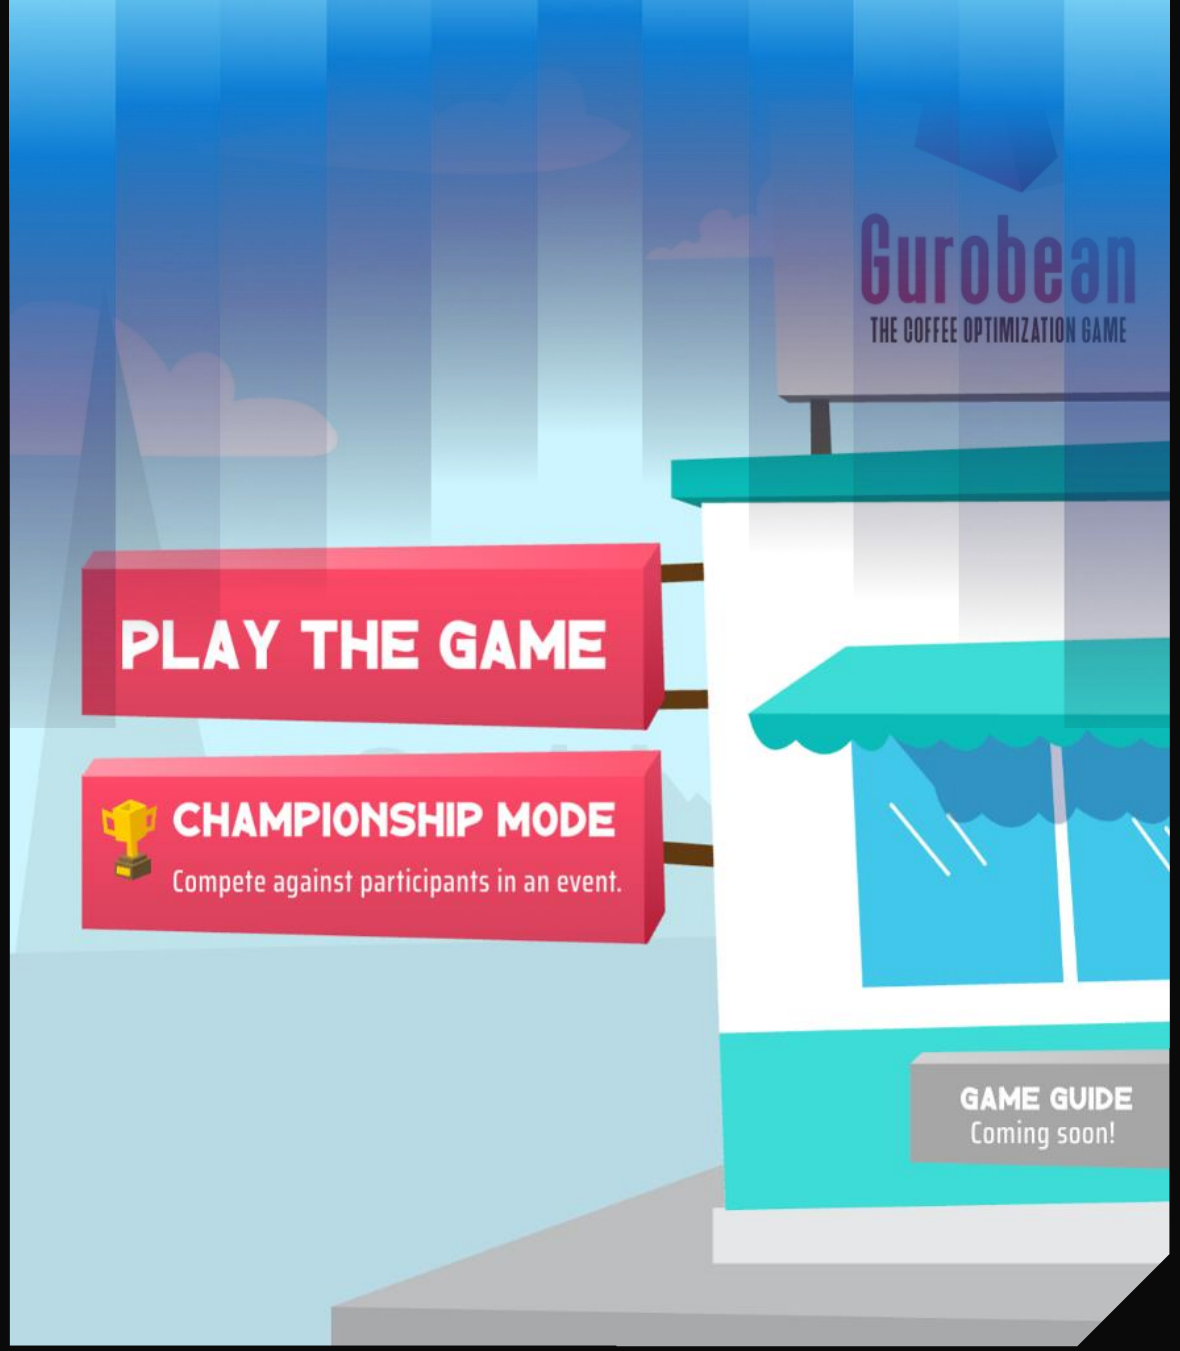

Gurobi

BUSINESS

What Students Should Take Away

Business & analytics learning outcomes, the lasting points, no math required

- **A single outcome doesn’t measure a decision.** Because demand is random, the same choice gives different results.
- **Uncertainty makes performance inherently variable.** Two “good” days can follow the very same decision.
- **Feasible plans can perform very differently.** Meeting every constraint doesn’t make a plan good.
- **Systems with people and queues behave non-intuitively.** Small changes can overwhelm gut instinct.
- **Analytics supports decisions, it doesn’t eliminate risk.** Good decision-making is about managing uncertainty, not removing it.

BUSINESS

Mathematical Optimization

Every round of Gurobean is secretly one of these

maximize / minimize	**a performance metric**	“objective function”
subject to	**restrictions on decisions**	“constraints”
by choosing	**quantities we control**	“decision variables”

For the Gurobean coffee shop:

maximize	5-day profit
subject to	beans, water, and queue/service capacity
by choosing	brew rates, markup, and barista service rate

BUSINESS

Feasibility ≠ Success

Meeting every constraint is only the entry ticket.

- When the sliders turn red, the plan is infeasible, it breaks a resource or capacity limit. Students quickly learn to get back into the green.
- But among all the feasible plans, performance still varies enormously.
- **“Is this allowed?” and “Is this good?” are different questions.**

Ask the room: two feasible plans, wildly different profit, what separates a good decision from a merely allowed one?

BUSINESS

Trade-offs & Diminishing Returns

Every decision pulls in more than one direction, and more is not always better.

- Brew too little and customers leave empty-handed; brew too much and coffee becomes waste.
- Serve too slowly and the queue grows; raise prices and margins improve but demand falls.
- The first extra unit of brewing or capacity helps a lot; later ones help less as the system saturates.
- **Good decisions balance competing effects, there is rarely a single dial to turn.**

Activity: have students find the point where pushing a slider further stops helping, or starts hurting.

BUSINESS

The Best Choice Is a Moving Target

The shape of the profit curve changes from round to round and scenario to scenario.

- As the round, the ingredients on hand, and the prices change, the profit curve reshapes.
- Three things move: where the peak sits (the best decision), how high it reaches (the best 5-day profit), and how sharp it is.
- A flat peak is forgiving, roughly right is good enough. A sharp peak is unforgiving, small misses cost real money.
- A binding resource limit can cut the curve short, putting the best choice at the constraint, not the natural peak.
- **So the best decision must be re-found as conditions change, not memorized.**

Activity: put two Profit Monitor curves side by side and ask where the sweet spot is now, and how costly a small miss would be.

BUSINESS

Common Misconceptions to Address

Myth: **"The best decision should always win."**
Reality: Not in a random system, a strong plan can lose on a single run and still be right on average.

Myth: **"More production always means more profit."**
Reality: Over-brewing wastes money; profit peaks and then falls. More is not always better.

Myth: **"If it worked once, it's a good policy."**
Reality: One good day is one lucky draw. Judge a decision by its distribution of outcomes, not one result.

Gurobi

Play Gurobean

www.gurobean.com