

Running the Cafeteria on Data: A Measurement, Forecasting and Production System for a 4-Station University Dining Hall

What the industry knows about cafeteria demand and waste, what our operation does not yet know about itself, and a staged plan to turn receipts, scales and a calendar into fewer rushes, fewer pans in the bin and a smarter order sheet.

1,600

stated diners per day (350 breakfast
· 730 lunch · 520 dinner)

4

service stations, staffed, with an
existing recipe base

0

data points currently captured on
sales mix, production or waste

Prepared for: **UH Dining leadership team** · Prepared by: **Sonny Steele** · Date: **September 5, 2026**

Status: **Research and proposal — no production or waste data has been collected yet.** Site facts confirmed by staff on September 5, 2026: Sodexo-operated UH Mānoa unit, all-you-care-to-eat entry, every diner card-scanned at front-desk check-in, no station point-of-sale, no costed recipes on site. Every other figure is either cited from an external source, computed from the three headcount averages above, or explicitly labeled as an assumption. Nothing here should be quoted as a measurement of this cafeteria until Phase 0 (Section 12) has run.

1. Executive summary

The honest starting point. We serve roughly 1,600 meals a day across four stations and cannot say, for any single dish, how many portions were cooked, eaten, or thrown away. That is a measurement problem, not a data-science problem. Until we log production and waste per item, per station, per meal, for at least four weeks, no model, dashboard or vendor can tell us anything we do not already guess.

What the industry says the prize is. Kitchens that measure kitchen (pre-consumer) waste with a scale and a reason code routinely cut it by 30–50% within a semester, and that waste typically sits at 4–10% of food purchases by value before tracking starts. Post-consumer plate waste in college dining runs 2.2–3.3 oz per diner per meal. Institutions that forecast per menu item with even simple smoothing models out-perform "same as last week" guessing, and a 2025 study on cafeteria turnstile data hit 8% daily forecast error using gradient-boosted trees with calendar features.

\$57k–\$143k

kitchen waste per year at the cited 4–10% band, on an *assumed* \$1.43M annual food purchase base

ASSUMPTION

\$23k–\$57k

recoverable per year at a 40% reduction, the mid-range of published tracking results

ASSUMPTION

49k–74k lb

plate waste per year at 2.18–3.3 oz per meal, published college benchmarks applied to our 358k meals

What we recommend, in order.

1. **Week 0: ask Sodexo for what already exists.** Three asks to the GM: the Transact door-swipe export, WasteWatch (Leanpath) enrollment status, and DRIVE recipe costs.
2. **Weeks 1–4: measure before you model.** A production record per recipe per station (forecast, cooked, leftover, disposition) and every pan that hits the bin weighed with a reason code. Paper plus a bench scale; one spreadsheet, entered the same day.
3. **Weeks 5–12: descriptive analytics.** Popularity index per item by weekday, waste by reason and station, stockouts per service, waste cost per meal. This alone says what we over-cook and what we run out of.
4. **Spring 2027: forecast and optimize.** Diner-count forecast (calendar-aware) × item preference = portions to prepare; then a newsvendor rule that sets a different safety margin for a \$3 seafood portion than for a \$0.15 scoop of rice. Batch plan: cook ~75% up front, hold the rest ready to fire.
5. **Then ordering.** Par levels and reorder points driven by forecast usage and Hawaii's barge lead times.

Three things leadership should not do

1. Do not buy tracking software before four weeks of paper logging proves the team will log. Compliance, not technology, is where these programs die.
2. Do not treat the 350 / 730 / 520 averages as fixed. They are real front-desk counts, which is good, but the period they cover is unknown; recompute them by weekday from the check-in records for the last full semester.
3. Do not put the "printed receipts" plan at the center. In an all-you-care-to-eat hall a receipt records a body at the door, not a dish at a station. The production record per station is the item-level data.

Contents

- 1 Executive summary
- 2 What we know, what we assume, what we do not know
- 3 How university cafeterias actually run: the evidence
- 4 The measurement system: five data streams and a data model
- 5 The analytics: from popularity index to forecast to production quantity
- 6 Calendar, holidays and the semester as a demand regime
- 7 Ordering and inventory: par levels, lead times, Hawaii logistics
- 8 Operational changes that do not need a model
- 9 Tooling options and cost tiers
- 10 KPIs and the leadership dashboard
- 11 Brutal truths, risks and failure modes
- 12 90-day plan and decisions required
- 13 Sources
- A Appendix: production record and waste log templates
- B Appendix: prototype screens (Mānoa Dining Ops)

Evidence tags used throughout. **CITED** external source listed in Section 13. **ASSUMPTION** a number we chose so the arithmetic could run; replace it with our own data. **UNVERIFIED** a claim about our operation that has not been checked on site. **CONFIRMED** stated by site staff (Sonny Steele) on September 5, 2026; not yet backed by an export.

2. What we know, what we assume, what we do not know

Item	Status	Why it matters
Diners per meal: 350 / 730 / 520	CONFIRMED Digital check-in: every diner's card is scanned at the front desk. Averaging period unknown.	Bodies are known; dishes are not. Card scans run on the Mānoa One Card system (Transact), so every entry already carries a timestamp. The 15-minute arrival curve, weekday profile and holiday effects are sitting in that system waiting for an export.
Four stations, existing recipes, "a bunch of employees"	UNVERIFIED Station names, menus, staffing per shift not documented.	The production record is per station per recipe. We need the recipe list with portion sizes and ingredient costs.
Service days per year	ASSUMPTION 224 (two 16-week semesters × 7 days). Real number depends on weekend and summer service.	Sets annual meal volume (358,400) and every dollar figure in Section 1.
Raw food cost per meal	ASSUMPTION \$4.00. Mainland residential dining commonly budgets \$2.50–3.50 raw food; Hawaii imports 85–90% of its food and grocery prices run ~57% above the mainland, so we padded it. Sodexo contract pricing may be lower.	Sets the \$1.43M purchase base. Sodexo's DRIVE platform reports the unit's actual food cost; ask the GM for one semester's number.
Service model (all-you-care-to-eat swipe vs. pay-per-item retail)	CONFIRMED All-you-care-to-eat, one swipe at the door via Mānoa One Card (Transact).	Item-level sales do not exist and will not. Consumption per dish must be inferred as prepared – leftover from the station production record. That record is the backbone of the whole system.
Operator (self-op vs. Sodexo contract)	CONFIRMED Sodexo-run, UH Mānoa.	Sodexo's standard waste program is WasteWatch (Leanpath, 2,000+ sites) and its recipe/purchasing platform is DRIVE. Whether this unit is enrolled in WasteWatch is UNVERIFIED .
Fall 2026 academic calendar	CITED UH Mānoa registrar.	Holidays, recess and finals are the largest predictable demand shocks.
Station POS and costed recipes	CONFIRMED Neither exists at the site.	DRIVE may still hold per-portion costs centrally; ask. Otherwise costing the top 40 recipes by hand is a Phase 0 task.
Current waste volume, stockouts, popularity	Unknown. Nothing is tracked.	This is the gap the whole plan closes.

Three asks for the Sodexo general manager before anything is built

1. **Transact swipe export** for the last full semester: one row per swipe with timestamp and location. This is stream 1 of Section 4, already collected.

2. **WasteWatch status:** is this unit enrolled in Sodexo's WasteWatch powered by Leanpath? If yes, the tracker, reason codes and reports in Section 9 already exist and the job is to use them. If no, ask what enrollment requires; Sodexo has committed publicly to deploying it across thousands of sites.

3. **DRIVE recipe costs:** per-portion food cost for the recipes on the four stations. Without it the scale reports pounds and leadership hears nothing.

Whether the unit has WasteWatch or DRIVE access is **UNVERIFIED** ; that Sodexo operates it is

CONFIRMED .

Reading the rest of this document

Sections 3, 5, 6 and 7 hold up regardless of which assumptions above turn out wrong: they describe methods and published results. Section 1's dollar figures and Section 10's targets are placeholders that must be re-run with our own numbers after Phase 0. The simulation script in the project repository re-computes every chart from a single set of inputs, so the swap takes minutes.

3. How university cafeterias actually run: the evidence

3.1 How much gets wasted, and where

Finding	Number	Source / context
College plate waste per diner per meal (20 Bon Appétit sites, 2018–19)	2.18 oz	More than twice the corporate-cafeteria rate (1.03 oz). ≈112 lb per student per school year. CITED
Plate waste per meal, U. Illinois residential halls (2016 audit)	3.3 oz	≈14,875 lb per week across six halls. CITED
Share of selected food wasted, round plates → oval plates (Illinois, 1,200+ observations)	15.8% → 11.8%	~0.7 oz less waste per plate from plate shape alone. CITED
Kitchen (pre-consumer) waste as share of food purchases, by value	4–10%	Leanpath college data; post-consumer is often equal, so 8–15% total. CITED
Overproduction's share of all foodservice food waste	23%	ReFED Insights Engine, cited in the MCURC/ReFED "Repurpose with a Purpose" report. CITED
Return on food-waste reduction investment	\$14 per \$1	Cited in the same MCURC/ReFED report (originates from a Champions 12.3 / WRI business-case study). CITED

3.2 What measurement programs achieve

Program	Result	Notes
Sodexo + Leanpath pilot, 8 campuses, 8 weeks (2010)	≈ –33%	Kitchen waste cut by about one third; staff logged each discard with a reason at a tracking station; each site ran a "Stop Waste Action Team." CITED
Notre Dame Campus Dining with Leanpath	–30%	Pre-consumer waste; 17.9 tons of food saved since Nov 2015. CITED
Arizona State (post-consumer tracking)	–40%	Plate waste. CITED
Leanpath's general claim for daily digital pre-consumer tracking	–50% or more	Vendor claim; treat as an upper bound, not a plan number. CITED
Sodexo WasteWatch program average	–50%	Vendor/operator claim across 2,000+ sites. CITED

3.3 What operators actually do (survey of 57 U.S. university foodservice operations, 2019–20)

- 91% forecast demand to prevent overproduction; 88% prepare smaller batches; 86% are trayless; 84% adjust menu planning; 81% reuse leftovers in other dishes.
- 77% track food waste, but inconsistently: 47% at pre-consumer, 40% at service, 32% at post-consumer. The authors' conclusion: "food waste is not measured in a consistent way in universities across the U.S."

- Top barriers were labor (44% for donation, 42% for composting), liability concerns, and infrastructure.

CITED

Read-through for us: we are behind the 77% who track anything, but the practices that matter most (forecasting, smaller batches, repurposing) are procedural, not technological. We can adopt them in week one.

3.4 What forecasting achieves

- **Classic foodservice research (Ohio State, JADA):** diner count forecast with de-seasonalized exponential smoothing, then menu-item demand = count × a "preference statistic" per item. Simple smoothing and moving-average models beat the naive "same as last time" method at every site studied; exponential smoothing won at 4 of 7 sites, moving averages at the rest. CITED
- **Modern ML (Aydın, Balcıoğlu & Sezen, 2025):** two months of institutional cafeteria turnstile data. XGBoost with time features reached 8.3% daily MAPE; an LSTM handled 15-minute intervals at 11.6% MAPE; a weighted ensemble improved 4.7% on the best single model. Most important features: recent history, weekly cycle, academic calendar. Estimated waste reduction from better forecasts: 6.2%. Caveat: two months is a thin dataset. CITED
- **Other catering ML studies** using day-of-week, holiday, class-day and weather features report waste reductions from 14% to 52% with random-forest and LSTM models; the wide range says results depend heavily on data quality. CITED

Take-away

The winning recipe has not changed in 35 years: forecast the *bodies* with a seasonal model, forecast the *plates* with a per-item share, and re-fit weekly. Fancy models help at the margin only after the data exists. Our first forecast will be a spreadsheet; our tenth might be gradient boosting.

4. The measurement system: five data streams and a data model

The design goal is **ten seconds per log entry** and **one owner per stream**. Anything slower gets skipped on a busy Tuesday, and skipped logs are worse than no logs because they bias every average downward.

Stream	What is captured	How (Phase 0)	Owner / cadence
1. Headcount	Diners entering, per 15-minute bucket, per meal.	Already captured: every diner's card is scanned at the front desk (Transact). Request the per-scan export with timestamps; no new logging needed.	Front desk (already happening); manager pulls the export weekly.
2. Production record	Per recipe, per station, per meal: forecast portions, portions actually prepared (by batch, with time fired), portions left at close, disposition (reused / discarded / donated).	One paper sheet per station per meal (template in Appendix A). Photograph it at close; enter in the sheet same day.	Station lead. Every service.
3. Pre-consumer waste log	Every pan or container discarded in the kitchen: item, weight, reason code (overproduction / trim / spoilage / expired / burnt / other), station, stage (prep vs. service).	Bench scale by the waste bin with a laminated log. Tare the pan.	Whoever discards. Continuous. Reviewed daily by the chef.
4. Post-consumer plate waste	Total plate-scrappings weight for a sampled meal; optionally sorted by category.	Weigh the dish-room bin(s) at the end of two meals per week. Divide by headcount.	Dish room lead. Twice weekly, rotating meals.
5. Context calendar	Class day flag, holiday, recess, finals, campus events, weather, menu cycle day, any service disruption.	One row per date, filled a semester ahead from the registrar calendar; events and disruptions added as they happen.	Manager. Weekly.

4.1 What "printed receipts" can and cannot give us in an all-you-care-to-eat hall

- The swipe at the door is the only transaction. A printed receipt is a paper copy of a Transact row: who entered and when. It says nothing about which station they visited or what they took. There is no item-level sales data to print, and no amount of receipt printing will create it.
- What the swipe data **is** good for: the 15-minute arrival curve per meal, headcount by weekday, holiday and finals effects, and (with the plan type on the row) the resident vs. commuter mix. That is the entire diner-forecast side of Section 5 and it is already collected.

- Item consumption therefore has to be inferred as *prepared – leftover at close* per recipe per station, which is why stream 2 (the production record) is the backbone of the whole system and the one thing the team must start doing by hand.
- If Sodexo's WasteWatch tracker is present, stream 3 is already a tablet-and-scale workflow with reason codes; the production record still has to be added, because WasteWatch measures what was thrown away, not what was cooked or served.

4.2 Data model (six tables)

Column types are illustrative; this runs identically as six tabs in Google Sheets or six tables in Postgres.

service	date, meal, station, bucket_start(15min), diners, notes
production	date, meal, station, recipe_id, forecast_portions, batch_no, portions_prepared, time_fired, portions_left_at_close, disposition(reuse discard donate), leftover_lb
waste_log	timestamp, station, recipe_id_or_item, stage(prepare service plate), reason_code, weight_lb, logged_by
plate_sample	date, meal, diners_in_sample, total_plate_waste_lb, category_breakdown(optional)
calendar	date, is_class_day, holiday_name, recess_flag, finals_flag, event, weather_summary, menu_cycle_day, disruption_note
recipe	recipe_id, name, station, portion_oz, food_cost_per_portion, shelf_life_hours, repurposable(Y/N), prep_lead_minutes, batch_size_portions

4.3 The one derived table everything hangs off


Each night the sheet computes, per recipe and meal: $\text{served} \approx \text{prepared} - \text{left_at_close}$, $\text{take_rate} = \text{served} / \text{diners}$, waste_lb , $\text{waste_cost} = \text{waste_lb} \div (\text{portion_oz}/16) \times \text{cost_per_portion}$, and $\text{stockout_flag} = 1$ if the station reported running out before the meal ended. That table is the input to every chart and model in Section 5.

Recipes must be costed or the scale is useless to leadership


A scale reports pounds. Leadership decides in dollars. The bridge is a food cost per portion for every recipe on the four stations. If that costing does not exist, building it (ingredient quantities $\times$ current invoice prices) is a Phase 0 task and it is bigger than it sounds: expect one to two days for a 60–100 recipe base.

4.4 A student's lunch and the kitchen's day: where each data point comes from

Nothing about the student's experience changes. What changes is that each step now leaves a record. The first step is already captured today; the other four are the new work.



How the pieces connect



5. The analytics: from popularity index to forecast to production quantity

5.1 Phase 1 — descriptive (weeks 5–12): questions we can answer with a spreadsheet

Question	Computation
Which dishes are popular, by day of week?	Popularity index = portions served ÷ diners at that meal, averaged per weekday. Rank per station. A dish at 0.35 on Tuesday and 0.20 on Friday gets two different production targets.
What do we over-cook?	Leftover portions ÷ prepared portions, per recipe; rank by leftover <i>dollars</i> , not pounds. The top ten recipes will typically account for most of the waste cost.
What do we run out of, and when?	Stockout flags per recipe × time of stockout vs. the arrival curve. Running out at 12:40 with 300 diners still coming is a batch-timing failure, not a total-quantity failure.
Which station is the bottleneck?	Diners per 15 min vs. station throughput (portions plated per minute, observed twice). Queue forms when arrivals exceed service rate; rebalance staffing or pre-plate.
What does waste cost per meal?	Σ waste_cost ÷ diners, weekly. This is the headline KPI.
How wrong are our current guesses?	MAPE of forecast_portions vs. served, per recipe. This baseline is what every later model must beat.




Figure 1a. The single most useful table in the building: average diners by weekday and meal, from card scans. Tuesday lunch is twice Saturday lunch. Cooking the same amount on both days is where today's waste and today's rush both come from. *Simulated values built from the stated averages; the real table comes straight out of the Transact export.*




Figure 1b. What the waste log produces after four weeks: the ten dishes that carry most of the waste bill, ranked by dollars, not pounds. Expensive proteins rise to the top even when their leftover weight is small. *Simulated with a 21-dish sample menu; the real list is whatever the scale says.*




Figure 1c. The "we rush" map: share of dishes that ran out before the meal ended, by station, meal and weekday, under today's cook-the-same-amount practice. Dark cells cluster on Monday–Wednesday and on the small weekend dinner crew. *Simulated.*




Figure 1. The day-of-week profile we expect (Mon–Thu strong, Friday soft, weekends at roughly half). These multipliers are assumptions chosen to be typical of residential dining; Phase 0 replaces them with measured values. Dashed line = the stated average.

5.2 Phase 2 — forecasting (Spring 2027): bodies first, then plates

Step 1: forecast diners per meal. Start with weekly-seasonal exponential smoothing (Holt-Winters, period 7) with calendar overrides for holidays, recess and finals. Once we have a full semester, fit a gradient-boosted model (XGBoost or LightGBM) with features: same-weekday lag 1 and 2, 7-day mean, weekday, week-of-semester, class-day flag, holiday/recess/finals flags, event flag, menu-cycle day, rain flag. The 2025 turnstile study's 8% daily MAPE is a reasonable target; a spreadsheet Holt-Winters should reach 12–15%.




Figure 2. What the headcount view will look like once data exists. **The series is synthetic**, generated from a 730-diner lunch average with the Figure 1 weekday shape, real Fall 2026 holiday dates, a first-week bump and a finals dip. The dashed line is a plain weekly-seasonal smoothing forecast that learns as it goes; on this synthetic series it scores 13.8% MAPE against 18.1% for "same weekday last week." Holiday dips are where a calendar-aware model earns its keep.

Step 2: forecast portions per recipe.

$$\text{Portions}_{\text{item, meal, date}} = \text{Forecast diners}_{\text{meal, date}} \times \text{Take rate}_{\text{item, weekday}} \times \text{Competition adjustment}$$

The take rate is the smoothed popularity index for that weekday. The competition adjustment (start at 1.0) captures the fact that a dish's share depends on what else is on the line that day; after a semester of data it becomes a small regression on the co-offered menu. This is exactly the "count $\times$ preference statistic" structure validated in the Ohio State research.

5.3 From forecast to production quantity: the newsvendor rule

A forecast is a mean. We never cook the mean; we cook the mean plus a margin, and the right margin depends on the cost of being wrong in each direction:

$$\text{Critical fractile CF} = C_u \div (C_u + C_o) \quad Q^* = \text{the CF-th percentile of forecast demand}$$

- **C_o (overage)** = food cost of one leftover portion minus what it is worth if repurposed. Rice: pennies. Grilled mahi: dollars.
- **C_u (underage)** = cost of one portion short: a rushed batch (labor at overtime pace, energy), a disappointed diner, and pressure on the next station. We suggest \$2.00 per portion to start and let the chef tune it. **ASSUMPTION**




Figure 3. Three dishes with the same 200-portion mean demand ($\sigma = 30$) and very different optimal production. Cheap, repurposable rice should be cooked at the 93rd percentile (244 portions): running out is far worse than leftovers. Perishable seafood should be cooked *below* the mean (192 portions) because every leftover portion burns \$3. The repurposable entrée lands at 217. Today all three are probably cooked "about the same as last week."

5.4 Why forecast error is money in both directions




Figure 4. For a single entrée with a 182-portion mean (730 diners $\times$ 25% take rate) cooked at a 71% critical fractile, moving from 30% forecast error to 8% cuts expected leftovers from about 50 portions to 13 per lunch, and shortfalls from 13 to 3. Multiply by four stations, three meals, 224 days.

5.5 Batch plan: the fix for "rushing"

The rush the team feels is rarely a total-quantity error. It is cooking everything before service and discovering at 12:30 that the line is empty. The standard practice (88% of surveyed operators prepare smaller batches) is:

1. Cook **~75% of Q^*** before service, staged in the holding cabinet.
2. Hold the remaining **~25% as ready-to-fire** components (par-cooked, portioned, in the reach-in).
3. A **fire trigger** per station: when the arrival curve says 30% of diners are still to come and the hot well is below one-third, fire the next batch. The 15-minute headcount stream is what makes this a rule instead of a guess.
4. Final batch sized to the forecast of remaining diners, not to a full pan.

This one change usually cuts both the rush and the end-of-service leftovers, before any model is built.




Figure 4a. What the card-scan timestamps give us for free: the arrival curve. The dashed line marks the moment 70% of expected diners are in, which is the fire trigger for the last batch. Today nobody on the line can see this; it is sitting in Transact. *Simulated shape.*

5.6 Phase 3 — what becomes possible after a full year of data

- **Menu engineering:** popularity × contribution margin per recipe → promote "stars", re-price or re-portion "puzzles", drop "dogs". Requires costed recipes.
- **Menu-mix effects:** which pairings cannibalize each other across stations (taco day kills the grill).
- **Station load balancing:** queueing model (arrival rate vs. service rate per station) to shift staff across the meal rather than across the day.
- **Anomaly alerts:** a day whose waste or stockouts fall outside the expected band triggers a next-day review, automatically.
- **Semester-over-semester learning:** the Fall 2026 calendar effects become the priors for Fall 2027.

5.7 What a full-semester simulation says the plan is worth

To size the prize before any data exists, we simulated Fall 2026 day by day: 21 dishes across the four stations, every meal, the real registrar calendar, the stated headcount averages. Weeks 1–4 model today's practice (the same quantity every weekday, a smaller fixed quantity on weekends, little repurposing). From week 5 the plan in this section is switched on. **This is a model of the target effect, not a measurement.** The two things it is good for are showing what the dashboard will look like and checking that the plan's numbers sit inside the published 30–50% band rather than above it.

\$0.42 → \$0.29

kitchen waste per meal, weeks 1–4 vs. weeks 7–13 (–31%)

SIMULATED

5.0 → 1.1

dishes that ran out per meal service (–79%) **SIMULATED**

≈ \$45,875

kitchen-waste dollars per year at 224 service days, on ≈ \$149,453 at stake **SIMULATED**




Figure 4b. The two numbers leadership will watch, week by week. The Labor Day spike is the flat policy cooking a full Monday on a half-empty day; the plan's calendar override removes it. Both curves fall at the same time, which is the point: this is not a trade of one problem for the other. *Simulated.*




Figure 4c. Where the food goes. With hold-and-fire, part of the safety margin is never cooked; with a named day-two use, part of what is cooked comes back as tomorrow's special. The discarded share is what the scale weighs. *Simulated.*

How to read these numbers

The simulated baseline lands at about 10% of food purchases, the top of Leanpath's 4–10% band, because a flat cook-the-same-amount policy is the worst case the literature describes. If our real baseline is closer to 5%, the dollar savings roughly halve and the stockout reduction stays. The forecast-error figure (28% → 5.3%) compares "no forecast at all" with a calendar-aware model on clean synthetic data; on real data expect 8–12%.

6. Calendar, holidays and the semester as a demand regime

In Honolulu the weather is nearly constant, so the calendar, not the sky, is the dominant external driver. Fall 2026 at UH Mānoa (all ten UH campuses share the same start dates and holidays):




Figure 5. Fall 2026 from the UH Mānoa registrar: instruction Aug 24 – Dec 10; Labor Day Sep 7; Election Day Nov 3; Veterans Day Nov 11; Thanksgiving Nov 26 with a non-instructional day Nov 27; study period Dec 11–12; finals Dec 14–18. Spring 2027 runs Jan 11 – May 7 and adds Kūhiō Day (Mar 26), Good Friday (Apr 3) and Memorial Day (May 25). CITED

Regime	Expected effect (to be measured)	How it enters the model
Week 1 (move-in, orientation)	Above-average counts, unusual mix (families, new students exploring).	<code>week_of_semester = 1</code> flag; treat as its own regime, not as a forecast seed.
Steady state (weeks 2–13)	Strong weekly cycle; Friday soft; weekends low if residents leave campus. The same dish draws differently on different menu-cycle days.	Weekday seasonality; lag features; <code>menu_cycle_day</code> keys the take-rate once a full cycle is logged.
Single-day holidays	Sharp dip on the day; often a softer day before. Hawaii-only state holidays (Kūhiō, Kamehameha, Statehood) included.	<code>holiday</code> and <code>day_before_holiday</code> flags.
Thanksgiving recess	Multi-day drop; many residents fly home, so expect deeper dips than mainland benchmarks. Nov 25 partial; Nov 26–29 very low.	<code>recess</code> flag + <code>days-into-recess</code> .
Study period + finals (Dec 11–18)	Breakfast collapses, late-night demand rises, mix shifts to grab-and-go.	<code>finals</code> flag; consider different hours rather than a different forecast.
Campus events (games, concerts, career fair)	Localized spikes at one meal.	Manual <code>event</code> flag entered a week ahead.
Weather (rain, tropical storm)	Small for residential dining; larger for walk-in retail traffic.	Rain flag from NWS Honolulu; test whether it matters before keeping it.

7. Ordering and inventory: par levels, lead times, Hawaii logistics

Ordering is the second-order win. It only becomes data-driven once the production forecast exists, because the forecast is the demand signal the order sheet has been missing. Three rules:

7.1 Usage forecast per ingredient

$$\text{Ingredient usage}_{\text{week}} = \sum_{\text{recipes}} (\text{Forecast portions} \times \text{ingredient qty per portion}) \times (1 + \text{trim/yield loss})$$

Explode the recipe forecast into ingredients using the recipe file. Yield loss (peeling, trimming, cooking shrink) is exactly what stream 3's "trim" reason code measures.

7.2 Reorder point and safety stock

$$\text{Reorder point} = \text{mean daily usage} \times \text{lead time (days)} + z \times \sigma_{\text{daily usage}} \times \sqrt{\text{lead time}}$$

Use $z = 1.28$ (90% service) for shelf-stable and frozen items, $z = 0.85\text{--}1.0$ for fresh produce and proteins where overage costs more, and re-derive from the newsvendor logic where costs are known. Hawaii's lead-time reality: most dry goods and much protein arrive by barge on weekly cycles, so a missed order cannot be fixed with a next-day truck. That argues for higher safety stock on shelf-stable items and **lower** production margins on fresh items, which is the opposite of what a mainland playbook assumes.

7.3 Par levels and shelf life

- Set each item's par = reorder point + order quantity, and cap it by shelf life: never hold more than (shelf life $\div$ lead time) cycles of usage.
- Log "expired" and "spoilage" waste against the ingredient, not the recipe. Those codes are the audit trail that tells purchasing the par is too high.
- First-in-first-out with dated labels is non-negotiable; the waste log will show you the day it slips.

7.4 Inventory counting cadence

Weekly count of the top 40 ingredients by spend (they will be 80% of the dollars). Everything else monthly. Variance = expected usage – actual usage; large negative variance is theft, over-portioning, or an un-logged discard, and the production record tells you which.

8. Operational changes that do not need a model

Everything below is standard practice in the operations surveyed in Section 3.3 and can start the same week as data capture.

Practice	What to do	Evidence
Smaller batches, hold-and-fire	Section 5.5. Stop cooking to a full pan; cook to the forecast and fire the rest on a trigger.	88% of surveyed operators.
Production sheet discipline	No recipe is cooked without a forecast number written on the sheet first. The sheet is the forecast's diary.	91% forecast; Ohio State methodology.
Repurposing plan per recipe	Every entrée has a named "day-two" use (soup, grain bowl, fried rice, staff meal). Mark it in the recipe file. Repurposing sits one rung below prevention on the EPA Wasted Food Scale.	MCURC/ReFED sprint: ~\$20k food cost saved in one month across nine institutions.
Portion tools	Standard scoops and spoodles per recipe; the portion_oz in the recipe file has to be the tool actually on the line.	Portion optimization is the first-listed prevention strategy in the MCURC report.
Plate size and trayless service	If we still use trays, remove them (25–30% plate-waste reduction reported). If plates are large, test a smaller or oval plate on one line for two weeks and weigh the dish-room bin.	Illinois 15.8% → 11.8%; Leanpath trayless data.
Cooling and carryover food safety	Leftovers destined for reuse must hit 135→70°F in 2 h and 70→41°F in 4 h with a logged cooling record. Reuse without the log is a health-code risk, not a savings.	Standard HACCP cooling limits.
Waste bin visibility	Put the scale and log next to the bin, in sight of the line. The act of weighing changes behavior before any report is read (the Hawthorne effect is a feature here).	Every Leanpath case study.
Weekly 15-minute waste huddle	Chef + station leads read the top five waste items by dollar and pick one fix each. Sodexo's pilot called this the SWAT team.	Sodexo/Leanpath 2010 pilot.

9. Tooling options and cost tiers

Tier	Stack	Approx. cost	When
0 — Paper + sheet	Laminated logs, one bench scale per kitchen (30–60 lb capacity, 0.5 oz resolution), a Google Sheet with the six tabs in 4.2, charts in the same sheet.	Scale \$60–150; labor only otherwise. ASSUMPTION	Now. Weeks 1–4. Proves compliance before any purchase.
1 — Tablet forms	Google Forms or AppSheet on a \$150–300 tablet at the waste station and at each line; scale with a display or a Bluetooth scale; Looker Studio (free) dashboard on top of the sheet; the simulation script in this repo re-run weekly.	\$300–800 hardware; \$0 software. ASSUMPTION	Weeks 5–12, once paper works and entry time is the bottleneck.
2 — Sodexo WasteWatch (Leanpath)	Sodexo's own program: integrated scale + tablet, reason codes, automatic reports, corporate coaching. Deployed at 2,000+ Sodexo sites.	Operator program; normally funded inside the Sodexo contract, not by the site. Enrollment status here UNVERIFIED	Ask in Week 0. If enrolled, Tiers 0–1 collapse into "use it daily." If not, request enrollment before buying anything else.
Transact swipe export	Per-swipe CSV (timestamp, location, plan type) from the Mānoa One Card system that already gates the door.	\$0; a report request to the Sodexo GM or the UH card office.	Week 0. It is stream 1 in full and the diner-forecast dataset for Section 5.
DRIVE recipe costs	Per-portion food cost and ingredient quantities from Sodexo's menu/purchasing platform.	\$0 if the unit has access; otherwise 1–2 days of manual costing for the top 40 recipes.	Week 0–2. Required before any waste number can be reported in dollars.
Forecasting	Sheet formulas (Holt-Winters) → Python (statsmodels / Prophet / LightGBM) in the repo → optionally a scheduled job that writes tomorrow's production sheet.	\$0 software; analyst time.	Spring 2027.

Build vs. buy, brutally

Tier 2 vendors sell coaching and habit formation as much as hardware; their published results come with a dedicated program manager on site. If we cannot commit a named person to own the data daily, a \$10k tracker will produce the same nothing that a \$100 scale produces. Decide the owner first, the tier second.

10. KPIs and the leadership dashboard

KPI	Definition	Baseline	Target after 1 semester
Kitchen waste cost per meal	Σ pre-consumer waste \$ $\div$ diners, weekly	Unknown; measure	–30% vs. week-4 baseline
Kitchen waste as % of purchases	waste \$ $\div$ food purchases \$, monthly	Unknown; industry 4–10%	Below 4%
Plate waste per diner	dish-room bin lb $\div$ diners, sampled	Unknown; industry 2.2–3.3 oz	–15% (portion tools, plates)
Stockouts per service	count of recipe stockout flags before meal end	Unknown; "we rush"	≤ 1 per service, none in the first 75% of the meal
Forecast error, diners (MAPE)	$ \text{forecast} - \text{actual} \div \text{actual}$, per meal	Naive baseline, measured	$\leq 12\%$ spreadsheet; $\leq 8\%$ ML
Forecast error, top-20 recipes (MAPE)	same, per recipe	Measured in Phase 1	$\leq 20\%$
Leftover reuse rate	reused lb $\div$ (reused + discarded lb)	Unknown	$\geq 50\%$ of edible leftovers
Log compliance	production sheets completed $\div$ expected (4 stations $\times$ 3 meals)	0%	$\geq 95\%$ every week

Dashboard layout (one page): waste \$ per meal trend (weekly); top-10 waste recipes by \$ (this week vs. last); stockout heat-map (station $\times$ meal $\times$ weekday); forecast vs. actual diners (last 14 days); popularity index table by weekday for the current menu cycle; compliance gauge. The synthetic charts in this report are the prototypes for the first four panels.




Figure 6. The size of the prize under the assumed \$1.43M purchase base. The full-semester simulation in 5.7 lands at the top of this band (≈ \$149,453 at stake, ≈ \$45,875 recovered). Even the low case funds a Tier 1 stack in the first month of savings. Replace the base with the real purchasing ledger before this chart is shown as a projection.

11. Brutal truths, risks and failure modes

1. **We have no data, and every number in this report about our cafeteria is borrowed or assumed.** The report's value is the method and the plan, not the dollar figures. Anyone who quotes "\$57k" to a budget committee before Phase 0 is making it up.
2. **The averages are real counts over an unknown window.** 350 / 730 / 520 are front-desk check-in counts, which is the right source, but nobody knows the weeks they average. A semester average hides a 2× swing between Tuesday and Saturday and a 50% drop over Thanksgiving. Recompute them by weekday from the Transact export before they are used for anything.
3. **Compliance is the program.** Published 30–50% reductions came with a named owner, a daily review and a tracking station in the line of sight. Skipped logs bias every metric and quietly destroy the forecast. Design for ten seconds per entry and audit compliance weekly as its own KPI.
4. **Two months of data is not a model.** The best ML study we found trained on two months and its authors were careful about it. We need a full semester for the weekly cycle and calendar effects, and two semesters before Fall-vs-Spring differences are trustworthy. Do not promise an ML forecast before Spring 2027.
5. **Pre-consumer waste is ours; plate waste is the diner's.** Kitchen waste responds to production discipline within weeks. Plate waste responds to portion tools, plate size and signage over months. Sequence accordingly.
6. **The rush is a batch problem more than a forecast problem.** Hold-and-fire fixes most of it immediately. A perfect forecast cooked entirely before service still runs dry if the arrival curve is front-loaded.
7. **Uncosted recipes make the scale a sustainability toy.** Without cost per portion, the report to leadership is in pounds, and pounds do not move budgets.
8. **Governance is unclear and it matters.** Sodexo operates the unit; UH Student Housing is the client. Data, tooling and recipe systems belong to Sodexo. A parallel system built by site staff without the GM may duplicate WasteWatch or breach contract reporting norms. Decide in week 0 who the "leadership team" for this plan actually is: unit management, the Sodexo district manager, or the university client. The asks in Section 2 are the test.
9. **Food safety is a hard constraint on reuse.** Repurposing without cooling logs turns a savings program into a health-department problem.
10. **Hawthorne effect cuts both ways.** Waste will drop the week the scale appears, then creep back. The baseline must be measured in weeks 2–4, not week 1, and the trend reviewed monthly for regression.
11. **Hawaii logistics invert the mainland playbook.** Weekly barge cycles mean stockouts of dry goods are expensive to fix and fresh items are expensive to over-buy. Safety stock rules must be set per category, not copied from a textbook.

12. 90-day plan and decisions required

When	Do	Done means
Week 0	Name the data owner. Put the three asks in Section 2 to the Sodexo GM in writing: Transact swipe export for last semester, WasteWatch enrollment status, DRIVE recipe costs. Pull one month of purchasing totals.	Owner named; the three answers in writing; purchase total known.
Weeks 1–2	Buy one bench scale. Print the Appendix A sheets. Cost the top 40 recipes (from DRIVE, or by hand). Build the six-tab sheet. Train station leads in a 20-minute stand-up. Start all five streams.	≥ 80% of expected sheets returned in week 2.
Weeks 3–4	Baseline period. Fix log friction. Start hold-and-fire on the two highest-waste stations. Weekly 15-minute waste huddle begins.	Baseline KPIs (Section 10) computed from weeks 2–4 data.
Weeks 5–8	Popularity index by weekday; top-10 waste by dollars; stockout map. First forecast written on production sheets from data (weekday average × diner forecast). Move entry to tablet forms if paper is the bottleneck.	Production sheets carry a data-derived forecast every service.
Weeks 9–12	Newsvendor margins per recipe class (cheap/repurposable/perishable). Reorder points for the top 40 ingredients. First leadership dashboard. Thanksgiving recess observed and logged as the first calendar shock.	Week-12 KPIs vs. baseline reported with the receipts (the sheet) attached.
Spring 2027	Holt-Winters → ML diner forecast; calendar effects fitted from Fall; menu engineering with costed recipes; decision on Tier 2 tooling with real compliance and savings data.	Forecast MAPE ≤ 12% on diners; waste \$ per meal –30% vs. baseline.

Decisions leadership must make

1. **Who owns the data daily?** A name, with 30 minutes a day protected for it.
2. **Who puts the three asks to the Sodexo GM, and by what date?** Swipe export, WasteWatch status, DRIVE recipe costs.
3. **Will the four stations run production sheets on every service starting week 1?** This is a scheduling and culture commitment, not a technical one.
4. **Budget authority for Tier 0/1** (under \$1,000) now, and a decision date for Tier 2 in Spring 2027 based on measured results.
5. **Recipe costing:** from DRIVE if the unit has access; otherwise who costs the top 40 by hand, by when.

13. Sources

1. Bon Appétit Management Co. — "Study Finds College Students Waste 112 Pounds of Food Per School Year" (20 sites, May 2018–Mar 2019; 2.18 oz/guest/meal). bamco.com
2. University of Illinois ACES / phys.org — "Smaller plates help reduce food waste in campus dining halls" (2021; 15.8% → 11.8%; 3.3 oz/meal 2016 audit). phys.org · aces.illinois.edu
3. Leanpath — "How Colleges are Preventing Food Waste Today" (pre-consumer 4–10% of purchases; 50%+ reductions; trayless 25–30%). blog.leanpath.com
4. University of Notre Dame Campus Sustainability — "LeanPath Program helps Campus Dining Reduce Food Waste by 30%". green.nd.edu
5. Leanpath — "How Arizona State University cut pre- AND post-consumer food waste" (40% plate waste reduction). blog.leanpath.com
6. Sodexo / PR Newswire — "Sodexo Campus Food Waste Tracking Shows Dramatic Reduction in Kitchen Waste" (Nov 17, 2010; 8 campuses; ~one-third reduction; SWAT teams). prnewswire.com
7. Sodexo Australia — "WasteWatch powered by Leanpath" (program average ~50% prevention; 2,000+ sites). au.sodexo.com
8. Menus of Change University Research Collaborative & ReFED — "Repurpose with a Purpose: Operational Research Project Report" (overproduction = 23% of foodservice waste; \$14 saved per \$1 invested; ~\$20k saved in one month across nine institutions; EPA Wasted Food Scale). refed.org (PDF)
9. Food Waste Management Practices and Barriers to Progress in U.S. University Foodservice (survey of 57 operations, 2019–20). [PMC9180560](https://pubmed.ncbi.nlm.nih.gov/3180560/)
10. Food Choice and Waste in University Dining Commons — A Menus of Change University Research Collaborative Study (5 institutions, 296 subjects; AYCE and satisfaction effects). [PMC8000212](https://pubmed.ncbi.nlm.nih.gov/3000212/)
11. Aydin, Balcioğlu & Sezen (2025) — "Machine Learning Techniques for Cafeteria Demand Forecasting: An Institutional Case", OPUS Journal of Society Research (XGBoost 8.32% daily MAPE; LSTM 11.64% at 15-min; ensemble +4.7%; est. 6.2% waste reduction). dergipark.org.tr
12. "Forecasting in foodservice: model development, testing, and evaluation" and "Forecasting menu-item demand in foodservice operations", Journal of the American Dietetic Association (Ohio State North Commons, 1988 data; count × preference statistic; smoothing beats naive). jandonline.org · [PubMed 3351162](https://pubmed.ncbi.nlm.nih.gov/3351162/)
13. Miller, McCahon & Bloss (1991) — "Food Production Forecasting with Simple Time Series Models", Hospitality Research Journal (exponential smoothing best at 4 of 7 sites). [doi.org](https://doi.org/10.1016/0195-0687(91)90004-8)
14. "Machine learning models for short-term demand forecasting in food catering services: A solution to reduce food waste" (random forest / LSTM; 14–52% waste reduction range). [researchgate.net](https://www.researchgate.net/publication/351116212)
15. ReFED — 2025 U.S. Food Waste Report and Food Waste Pact 2025 Data Report (sector context; foodservice donation rate 3%). refed.org
16. MIT OpenCourseWare — "Newsvendor Inventory Problem" (critical fractile $CF = C_u / (C_u + C_o)$). [ocw.mit.edu](https://ocw.mit.edu/courses/15-095-newsvendor-inventory-problem/) (PDF)
17. University of Hawai'i at Mānoa, Office of the Registrar — Academic Calendar Fall 2026 and 2026–2027 calendar. [manoa.hawaii.edu](http://manoa.hawaii.edu/catalog) · catalog.manoa.hawaii.edu
18. UH Mānoa Student Housing Services — Residential Dining (Sodexo-operated; Gateway Café and Hale Aloha Café hours and stations). manoa.hawaii.edu/housing · uhm.sodexomyway.com
19. Hawaii food import share (85–90%) and grocery price premium (~57% above mainland): Beat of Hawaii and Living in Hawaii cost-of-living coverage (consumer, not institutional, pricing). beatofhawaii.com · livinginhawaii.com
20. Hechinger Report — average college three-meal plan ≈ \$4,500/yr, ≈ \$18.75/day (meal-plan revenue context). hechingerreport.org
21. HACCP cooling limits (135→70°F within 2 h; 70→41°F within the next 4 h) — standard FDA Food Code practice; templates e.g. Michigan Dept. of Education "Menu Production Records", Tazewell CPS "Production Log". michigan.gov (PDF)
22. Campus POS/reporting capability: CBORD Higher Education Retail & POS; Grubhub Onsite menu management. cbord.com · onsite.grubhub.com

23. UH Mānoa meal-plan swipe system: Mānoa One Card on Transact (eAccounts portal, Blackboard Transact app); swipes + Dining Dollars structure. eacct-hawaii-manoa-sp.transactcampus.com · uhm.sodexomyway.com/meal-plan/faq · khon2.com
24. Sodexo DRIVE platform (integrated menu, recipe, purchasing and inventory management with menu cost visibility). prnewswire.com · foodservicedirector.com
25. Sodexo UH Mānoa sustainability page (trayless dining at all locations; ~30% local product; no WasteWatch mention). uhm.sodexomyway.com/explore/sustainability
26. Sodexo WasteWatch powered by Leanpath global rollout (1,200+ then 3,000-site commitment). blog.leanpath.com · sodexo.com

Simulation and charts: `scripts/simulate.py` in the project repository (`~/Developer/GitHub/uh-cafeteria-analytics`). Inputs: the three stated headcount averages, the Fall 2026 registrar dates, and the assumptions listed in Section 2. All demand series are synthetic and labeled as such on the figures.

Appendix A. Production record and waste log templates

A.1 Station production record (one per station per meal)

Date	Meal	Station	Lead	Diners (from door tally)			Weather / event	
_____	B / L / D	_____	_____	_____			_____	
Recipe	Forecast portions	Batch 1 portions / time	Batch 2 / time	Batch 3 / time	Left at close	Ran out? (time)	Disposition (Reuse / Discard / Donate)	Leftover lb

Rule: the forecast column is written *before* cooking. "Left at close" is counted in portions with the station's scoop, then weighed. Photograph the sheet at close; enter it the same day.

A.2 Kitchen waste log (at the scale, next to the bin)

Time	Station	Item	Stage (Prep / Service / Plate)	Reason (Overproduced / Trim / Spoiled / Expired / Burnt / Dropped / Other)	Net lb (pan tared)	Initials
------	---------	------	--------------------------------	--	--------------------	----------

Rule: nothing edible goes in the bin un-weighed. Ten seconds per line. The chef reads the log every morning and the top five items by dollars are on the whiteboard by 10 a.m.


A.3 Door tally (headcount per 15 minutes)

Meal	:00–:15	:15–:30	:30–:45	:45–:00	Hour total
Hour 1					
Hour 2					
Hour 3					
Meal total					


If a swipe/POS export exists, this sheet is unnecessary; the export gives the same curve to the second.

Appendix B. Prototype screens: Mānoa Dining Ops

A clickable prototype of the dashboard was built from the same simulated dataset so that the plan can be shown rather than described. It is a static website with no backend; every panel reads the file produced by the simulation. Live version and source: github.com/sonnysteele23/uh-cafeteria-analytics.




B1. Landing screen: the four headline numbers and the simulated-data banner that never leaves the top of the page.



B2. The student journey and the kitchen's day, side by side.



B3. Live check-ins: the arrival curve replays a Tuesday lunch and the fire trigger lights at 70% of forecast.




B4. The Monday-morning dashboard: waste per meal, stockouts, forecast vs. actual, top waste dishes.

Where the food actually goes.


A scale next to the bin, a reason code, and a cost per portion turn pounds into a number leadership can act on.

Portions prepared per week: eaten, held, repurposed, discarded



Kitchen waste dollars by reason

Overproduction is the lever. Trim and spoilage are smaller and slower to move.



Kitchen waste dollars by station, per 4 weeks

Which line to stand at on Monday.



Waste log, last entries

WHAT THE SCALE STATION RECORDS						
TIME	STATION	ITEM	STAGE	REASON	NET LB	\$ BY
14:07	Greens & Simple Servings	Ahi Poke Bowl	Service	Overproduced	1.7	\$11 RL
14:05	Grill	Grilled Mahi Sandwich	Service	Overproduced	5.6	\$31 RA
10:42	Grill	Teriyaki Burger	Prep	Trim / prep	4.5	\$12 JT
14:03	Island Kitchen	Chicken Katsu Curry	Service	Overproduced	4.1	\$8 RL
09:15	Island Kitchen	Kalua Pork & Cabbage	Prep	Spoiled	2.5	\$5 RN
13:58	Wok & Global	Beef Broccoli Stir-Fry	Service	Overproduced	6.9	\$19 RA
12:31	Grill	Loco Moco	Service	Dropped	6.2	\$15 JT

Tuesday, October 20 lunch, generated the night before.

Forecast diners x each dish's Tuesday take-rate, then a margin sized by what it costs to be wrong. Cheap and repurposable dishes get a generous margin. Fresh fish does not.

Forecast diners		Dishes on the line		Cook before service		Hold ready-to-serve	
748		15		2,360 portions		786 portions	
STATION	DISH	CLASS	FORECAST	PERCENTILE	MAKE (Q+)	COOK NOW	HOLD DAY-TWO USE
Greens & Simple Servings	Garden Salad Bar (per plate)	PRODUCE	211	98th	213	175	58 — (oompost)
Greens & Simple Servings	Simple Servings Baked Chicken	ENTREE	108	98th	119	90	38 Chicken salad
Greens & Simple Servings	Ahi Poke Bowl	PERISHABLE	125	98th	118	88	29 — (never hold)
Grill	Teriyaki Burger	ENTREE	173	98th	178	134	45 —
Grill	Loco Moco	ENTREE	157	98th	171	128	43 Staff meal / hash
Grill	Grilled Mahi Sandwich	PERISHABLE	117	98th	112	84	28 — (look to order)
Island Kitchen	Steamed Rice	CHEAP	423	96th	543	487	136 Fried rice base
Island Kitchen	Fried Rice	CHEAP	226	92th	277	268	69 Next service
Island Kitchen	Spam Musubi	CHEAP	229	98th	257	193	64 —
Island Kitchen	Kalua Pork & Cabbage	ENTREE	222	98th	246	184	61 Kalua pork fried rice
Island Kitchen	Chicken Katsu Curry	ENTREE	215	98th	238	179	60 Katsu curry bowl special
Wok & Global	Chicken Adobo	ENTREE	170	98th	189	142	47 Adobo fried rice
Wok & Global	Beef Broccoli Stir-Fry	ENTREE	166	98th	180	135	45 Beef fried rice
Wok & Global	Saimin	ENTREE	144	98th	155	116	39 —
Wok & Global	Vegetable Curry (vegan)	ENTREE	112	98th	129	97	32 Curry soup

Percentile = the newsvendor critical fractile, cost-of-shortage + cost-of-leftover, floored at the 50th percentile for service level. "Hold" is pan-cooked and fried only if the card scans say the diners are coming.

B6. Popular days and tomorrow's production sheet: forecast, cost-aware make-quantity, cook-now and hold columns, day-two use.

B5. Food waste: where the food goes, by reason, by station, and the scale-station log.

Screens captured from the prototype on September 8, 2026. All values simulated.

Prepared September 5, 2026 (updated September 8) for the UH Dining leadership team. Research assembled from the public sources listed in Section 13; no proprietary or operator-internal data was used. Figures labeled SYNTHETIC or ASSUMPTION must not be quoted as measurements of this cafeteria.