HAPPYCHEF HOSPITALITY LIBRARY

FOR INDEPENDENT PIZZERIA OWNERS

The Multi-Location Pizzeria Playbook

How Independent Pizzeria Owners Grow From One Pizza Restaurant To Five Without Becoming A Chain

WRITTEN BY

Thibault Van de Sompele

HAPPYCHEF HOSPITALITY LIBRARY

The Multi-Location Pizzeria Playbook

How Independent Pizzeria Owners Grow From One Pizza Restaurant To Five Without Becoming A Chain

WRITTEN BY

Thibault Van de Sompele

Contents

Introduction

I have never made a pizza for money. I want to say that up front, because most books on this subject are written by people who ran a dough mixer for twenty years and then decided the rest of the industry needed their memoir. I come at this from a different angle. I build software that sits inside the daily operations of restaurants and hotels — the booking system, the floor plan, the kitchen ticket flow, the numbers that owners pull up at eleven at night when they are trying to work out why last Tuesday's labour cost looked so much worse than the Tuesday before. Over the years that software has sat inside hundreds of independent hospitality businesses across Europe, and a meaningful share of them make pizza.

What I noticed, watching from that seat, is that pizza is a strange business to scale. It looks simple from outside — flour, water, salt, yeast, a hot box, a topping list any customer can understand in four seconds. That simplicity is exactly why so many good cooks and so many good businesspeople open a second site convinced that what worked once will work twice, at roughly the same margin, with roughly the same effort. It rarely does. I have watched owners who ran a spectacular single site — queues out the door, five-star reviews, a dough they were quietly proud of — open a second location eighteen months later and watch the whole thing wobble. Not collapse, usually. Wobble. The dough at the new site is close but not right. The manager they hired is competent but not obsessive, and pizza rewards obsession in a way few other foods do. The delivery platforms that felt like a

nice add-on at one site become a third of revenue at two sites, at a commission that quietly resets the whole margin structure. And the owner, who used to be behind the peel every single service, is now driving between two locations, arriving at each one slightly too late to fix the thing that just went wrong.

This book exists because almost nothing written about restaurant growth is written for pizza specifically. The general management books will tell you to systemise, to hire well, to watch your numbers — all true, all useless without the specifics. Pizza has its own physics. Dough is a living culture with a clock running on it from the moment it is mixed, and that clock does not care about your expansion timeline. An oven has a hard number of pizzas it can produce per hour, and that number is your entire site's revenue ceiling whether you like it or not. Delivery platforms take a bite of margin so large that a badly configured delivery mix can make a full restaurant lose money on a busy Friday night. None of this shows up in a generic playbook about opening a second restaurant, because none of it is generic. It is pizza-specific, and it is exactly the terrain this book covers, chapter by chapter, with real numbers you can hold your own business up against.

I want to be honest about what this book is not. It is not a manual for building a national chain — there are other books, written by people who did exactly that, for that ambition. This book is for the owner of a genuinely good single pizzeria who is opening, or has just opened, a second and third location, and who wants five locations in five years that all still taste like the first one — not fifty locations in ten years that taste like nothing in particular.

That is a specific, achievable, and in my experience deeply underserved ambition. Most of the operators I have watched through this lens are not trying to become Domino's. They are trying to stay independent while getting bigger, which is a much harder needle to thread than either staying small or going fully corporate, because it means building just enough system to hold quality across distance without building so much system that you stop being the business your regulars fell in love with.

Every chapter in this book follows the same shape. I describe a situation you will recognise, because I have watched some version of it play out inside the businesses using our platform. I give you a named framework or calculation — something with a name you can say out loud to your management team, not just a vague piece of advice — with a worked example using invented but realistic numbers, always labelled as illustrative so you know to run your own figures rather than mine. And I end every chapter with something you can do this week, because a book about scaling a pizza business that does not change what you do on Monday morning has failed at its one job.

Twenty chapters, one business, one product. Let's get into it.

CHAPTER 1

Why The Second Pizzeria Fails More Often Than The First

Nobody opens a second pizzeria expecting it to be worse than the first. That is precisely why it usually is. The first site had a founder who tasted the sauce every morning, who noticed when the mozzarella supplier changed their moisture content, who could feel in their hands when the dough had proofed a half-hour too long. The second site has a manager who is trying their best with a recipe card and a phone number to call if something goes wrong. The gap between those two situations is where most second pizzerias quietly lose the thing that made the first one work.

I've watched this play out with a consistency that stopped surprising me years ago. An owner runs a first site for three, four, five years. It becomes genuinely excellent — not through a written system, but through thousands of small corrections the owner made in real time, mostly without noticing they were making them. They straighten a peel angle here, they tell someone to pull a pizza ten seconds early there, they taste the sauce and quietly add a pinch more salt before service without writing it down anywhere. None of it is documented because none of it needed to be — the person who knew it was there every single day. Then a second site opens forty minutes away, staffed by people who never watched any of those thousand corrections happen, working from a recipe card that captures

maybe a tenth of what actually made the first site work. The pizza that comes out is close. Close is the trap. Close reviews well for the first few weeks on residual goodwill from the brand name, and then it doesn't, and the owner is left wondering what went wrong when the honest answer is that nothing went wrong — nothing was ever specified precisely enough to go right.

I want to name the five failure modes I see most often, because recognising the pattern early is most of the fix.

The first is **recipe drift without detection**. The dough recipe on paper is identical across sites, but the actual dough is not, because paper recipes leave out everything that matters — ambient kitchen temperature, water temperature by season, how long the mixer actually runs before someone gets pulled away to answer the phone. Nobody notices for weeks because nobody is tasting both sites' pizza on the same day.

The second is **management by hope**. The owner hires a manager, tells them "just do what we do at the first site," and checks in by phone once a week. Hope is not a system. A manager who has never been told exactly what "the way we do it" means in numbers will invent their own version of it within a month, usually a slightly cheaper, slightly faster version, because that is what running a kitchen under pressure pushes you toward.

The third is **capacity mismatch**. The new site is chosen because the rent is good or the location looks busy, without anyone checking whether the oven, the prep space, or the delivery radius can actually support the volume the owner is

hoping for. A site that looks like a great real estate deal and a site that can produce four hundred good pizzas on a Saturday night are two different questions, and owners routinely answer only the first one.

The fourth is **channel shock**. The first site built its reputation on dine-in and walk-up regulars. The second site, in a new neighbourhood with no existing reputation, ends up leaning far more heavily on delivery platforms just to get any volume at all — and the owner discovers, usually around month three, that the margin structure of a platform-heavy site is nothing like the margin structure they are used to.

The fifth is **founder bandwidth collapse**. This is the quiet one. The owner does not stop caring, they just run out of hours. Two sites, two Friday nights, one person. Something has to give, and it is almost always quality control, because quality control is the task with the least visible short-term cost when it is skipped.

Here is a tool I give owners before they sign a lease on site two: **the Copyability Scorecard**. It asks a blunt question — is your first site actually a system, or is it just you? Score your first location from 0 to 4 on each of six dimensions: does a written dough specification exist that a new hire could follow without you present; does a written recipe and portion spec exist for every menu item; can your current best pizzaiolo train a new hire to an acceptable bake standard within two weeks without your personal involvement; do you have a written opening and closing checklist that isn't in your head; does your ordering and inventory process run without you physically checking stock; and can

someone else currently run a full Friday dinner service at your existing site while you are not there, with results you'd be comfortable a customer saw. Six dimensions, 0 to 4 each, 24 points total.

Illustratively: an owner scores themselves honestly and lands at 11 out of 24 — decent food knowledge transferred casually to staff, nothing written down, and a hard "no" on running a Friday service without them. That is not a business ready to copy. That is a talented individual with a following. The fix is not to cancel the second site — it's to spend eight to twelve weeks closing the gap on paper before you spend a single euro on a second lease, because every point you score above roughly 16 out of 24 corresponds, in what I've watched across dozens of second-site openings, to a meaningfully lower chance the second location needs an emergency rescue trip in its first ninety days.

Those eight to twelve weeks are not wasted time, even though they feel like delay when you're eager to sign a second lease. They are the period where you convert the thousand small unwritten corrections from the first paragraph of this chapter into documents, checklists and training routines a stranger could follow. Every week you spend writing the Dough DNA Sheet from the next chapter, filming a five-minute stretch-and-bake training video, or timing how long it actually takes a new hire to hit an acceptable bake standard, is a week that lowers the odds you'll be making an emergency ninety-minute drive to the second site on a Saturday night six weeks after opening. I have watched owners skip this step because momentum feels good and paperwork feels like the opposite of growth. It isn't. Paperwork,

done right, is what growth is actually made of in a business where quality lives in someone's hands and someone's timing.

This week: score your own first site honestly on the Copyability Scorecard, ideally with your head chef or best pizzaiolo scoring it independently and comparing notes with you, because the gap between your score and theirs is usually where the real problem is hiding.

CHAPTER 2

Dough Is The Business

Ask any pizzeria owner what they sell and most will say pizza. That's the wrong answer, or at least the answer one level too shallow. What you actually sell is a specific, repeatable dough performance, three times a day, under pressure, without fail. Everything else on the menu — the sauce, the cheese, the toppings — is decoration on top of a fermentation process that either works or doesn't. Get the dough right and a mediocre topping selection will still produce happy customers. Get the dough wrong and the finest San Marzano tomatoes and freshest buffalo mozzarella in the world will not save the plate.

This matters more, not less, once you have more than one site, because dough is a living thing on a clock, and clocks do not scale by copy-paste. A single-site owner manages one fermentation environment they know intimately — one kitchen's ambient temperature, one walk-in's actual running temperature versus its dial setting, one water supply's mineral content and seasonal temperature swing. A second site has all of those variables reset to unknown values, and dough is unforgiving about unknown values. Two hours of proofing at 22°C produces a meaningfully different dough than two hours at 19°C, and most kitchens have no idea their actual ambient temperature swings that much between a slow Tuesday and a packed Saturday with three ovens running.

The tool I put in front of every owner opening a second site is **the Dough DNA Sheet** — a single document that specifies dough not as a recipe but as a set of controlled variables and acceptable ranges, so that two sites making "the same dough" can actually be checked against each other rather than just trusted to be the same. It covers seven fields: flour type and protein percentage, hydration percentage, salt percentage, yeast type and percentage, mix time and final dough temperature target, bulk fermentation time and temperature, and final cold-proof time and temperature range. Every field gets a target and an acceptable tolerance band, not just a single number, because a spec with no tolerance is a spec nobody can actually hit under real kitchen conditions.

Illustratively, a Neapolitan-style dough spec might read: flour 00, protein 12.5%; hydration 62%, tolerance ±1%; salt 2.8%; fresh yeast 0.15%; mix to a final dough temperature of 24°C, tolerance ±1°C; bulk ferment 2 hours at room temperature, tolerance ±20 minutes; cold proof 20 to 24 hours at 4°C. That's not a poetic description of a family recipe. It's a document a new site's morning prep cook can be tested against with a thermometer and a scale, and a document you can hand to a manager and say "this is the standard, not a suggestion."

Two details make the Dough DNA Sheet actually work in a working kitchen rather than sitting laminated and ignored on a wall. First, every field needs an instrument, not a guess — a probe thermometer for dough temperature, a scale for hydration and salt rather than a jug and a pinch, a fridge thermometer that gets checked against the dial rather than trusted blindly, because

walk-in coolers drift from their displayed setting more often than owners assume, especially in an older second-hand unit bought cheap for a new site. Second, the sheet needs an owner-of-record at each site — one named person whose job includes checking the sheet against reality once a week, not "whoever's on shift," because a standard nobody is accountable for checking is a standard that exists only on the day you wrote it.

The second half of the tool is **the Fermentation Clock** — a simple schedule model that works backwards from your first scheduled bake of the day to fix every prior step's start time, so that dough readiness stops being a feeling and becomes an appointment. If your first pizza needs to leave the oven at 11:30am, and your cold proof needs 22 hours minimum, your mix needs to happen by roughly 1:30pm the day before, accounting for the two-hour bulk ferment and a buffer for staff realities like someone arriving ten minutes late. Write that backward schedule out for every daypart your site operates, post it by the mixer, and you have converted "the dough guy knows when to start" — a single point of failure — into a system that survives that person's day off.

Here's where the economics show up, and this is the part owners underestimate: a dough spec with tight tolerances directly reduces waste. A site running loose fermentation control typically throws away dough balls that over-proofed and went slack, or under-proofed and won't stretch properly, at a rate that, illustratively, runs 4–7% of daily dough production at an undisciplined site versus 1–2% at a site running a written spec with a posted clock. On a site mixing 300 dough balls a day at

roughly €0.55 in ingredient cost per ball, that gap between 6% and 1.5% waste is about €7.40 a day, or close to €2,700 a year — money that funds a genuinely meaningful chunk of a second site's opening costs, sitting inside a process most owners have never measured because it never occurred to them dough waste was a line item worth tracking separately from food cost in general.

Multiply that gap across three sites, each running its own fermentation control with its own degree of discipline, and the annual figure stops being a rounding error and starts being close to a junior manager's salary, sitting quietly inside a cost category most P&Ls bury inside a general "food cost" line where nobody ever isolates it. The owners who catch this are the ones who start weighing discarded dough balls for two weeks running at every site — not estimating, actually weighing what goes in the bin at close each night and logging it against that day's total mixed — because a number you can see on a clipboard changes behaviour in a kitchen faster than any policy memo ever will.

This week: write your Dough DNA Sheet for your current best-performing site's dough — every field, with tolerances, not just targets — and post your Fermentation Clock next to the mixer. If you can't fill in every field from memory without checking, that's the gap your second site is about to inherit.

DOUGH DNA SHEET — EXAMPLE SPECIFICATION

Parameter	Specification
Hydration	62–65%
Bulk ferment	2 hours, room temperature
Cold proof	24–72 hours, 3–4°C
Ball weight	250g ± 5g
Out-of-fridge rest before use	60–90 minutes
Same across every site, every day	**No local variation permitted**

CHAPTER 3

Commissary Or In-Store: The Production Decision

Somewhere between site two and site three, almost every pizzeria owner asks the same question: should we keep making dough at every location, or should we build one central kitchen and distribute it? It's a good question asked at the wrong time by most owners, because they ask it emotionally — "I'm tired of training dough guys at every site" — rather than mathematically. It's a real decision with a real break-even point, and getting it wrong in either direction costs you money for years.

The emotional version of the question usually surfaces after a bad week. A new hire at site three has been mixing dough for six weeks and still can't hit the hydration target consistently, a batch goes out under-proofed on a Saturday, and the owner, exhausted, decides on the spot that centralising production would solve everything. Sometimes it would. Often, at three sites, it wouldn't — it would trade one problem (training dough teams at multiple locations) for a different, more expensive one (a delivery van, a cold chain, and a single point of failure that now affects every site at once instead of just one). The decision deserves the same numeric treatment as everything else in this book, not a decision made on the worst night of the month.

In-store production keeps dough fresh, keeps every site's pizzaiolo skilled and invested in the product rather than just

baking someone else's puck, and avoids the cost and complexity of a delivery van doing daily runs. It also means you're managing fermentation quality at every single site independently, which is exactly the multiplication problem from the last chapter — more sites means more places the Dough DNA Sheet can quietly drift.

Commissary production centralises the skill and the quality control into one kitchen run by your best people, distributes a consistent product to every site, and frees up store-level staff and space for service rather than production. It also adds a transport cost, a transport risk (a late or failed delivery is now a same-day crisis at every site downstream of it), and a real question about whether dough that travelled two hours in a chilled van bakes identically to dough that never left the room it was mixed in.

The tool here is **the Commissary Break-Even** — a model that compares the fully-loaded cost of in-store production against centralised production plus distribution, as a function of site count, so you can find the number of locations at which the commissary actually starts saving you money rather than just feeling more organised.

Build it from four cost blocks. First, in-store production labour: the hours per week a site spends on dough production specifically — mixing, portioning, tray management — multiplied by the fully loaded hourly labour cost. Second, in-store waste: from Chapter 2, the ingredient cost lost to fermentation inconsistency, which tends to run higher per site when every site is training its own dough team from scratch. Third, commissary fixed cost: rent,

equipment, and the wages of a smaller, more skilled central production team, spread across however many sites you're serving. Fourth, distribution cost: the van, driver, fuel, and chilled packaging cost per delivery run, typically a daily or every-other-day cost per site depending on your production and shelf-life math.

Illustratively: a site doing its own dough spends roughly 14 labour hours a week on production at €16/hour fully loaded, plus about €95/week in avoidable fermentation waste — call it €319/week in production cost per site, in-store. A commissary serving that same site needs about 2.5 fewer labour hours locally (someone still portions and finals it), so in-store cost drops to roughly €135/week, but the site now needs a distribution run averaging €45/week, plus its share of the commissary's fixed cost — a central kitchen might run €2,400/week in rent, equipment lease, and skilled labour, which spread across four sites is €600/week per site. Total commissary-model cost per site: roughly €780/week. At four sites, commissary is more expensive than in-store, not less — the fixed cost is being carried by too few locations. Run the same numbers at eight sites and the fixed cost per site drops to €300/week, bringing the commissary total to about €480/week per site against an in-store cost of €319/week — now in-store is still cheaper, but the gap has narrowed sharply, and if your in-store waste is running high (undisciplined dough teams at each new site, which is common exactly when you're growing fastest), the crossover can happen as early as five or six sites.

The number that actually matters is not a universal "commissary at N sites" rule — it's your own break-even, and it moves with your waste rate, your local labour cost, and your commissary's fixed cost efficiency. Most owners I've watched get this right land the crossover somewhere between five and nine sites, almost never before four and rarely worth delaying past ten.

There's also a quality dimension the pure cost model doesn't capture, and you should weigh it deliberately rather than let the spreadsheet decide alone. Dough that travels two hours in a chilled van and finishes its cold proof in a second location's fridge is not automatically worse than dough mixed on-site — professional bakeries and pizza chains prove daily that it can be excellent — but it does add a step where something can go wrong that a single site never has to manage: a van breakdown, a temperature excursion during loading, a driver running late on the one morning your Saturday brunch rush needs dough by 10am sharp. Before you commit to a commissary model, run it in parallel with in-store production at one site for four to six weeks, blind-taste the results yourself against your Dough DNA Sheet standard, and only then decide whether the cost savings are worth the added logistics risk at your actual site count.

This week: build your own four cost blocks with your real numbers, even roughly, and calculate what your break-even site count actually is — then compare it to how many sites you're planning to have in three years, so you know today whether you're building toward an in-store forever business or a commissary you should start scouting space for now.

CHAPTER 4

Oven Throughput Is Your Capacity

Every pizzeria owner I've watched hit a growth wall eventually says some version of the same thing: "we were busy but we weren't making money that night." Nine times out of ten, the real story is that the dining room could seat more people than the oven could actually feed, and nobody had ever done the math to notice the mismatch before it cost them a fully booked, under-earning Saturday.

Picture the scene, because it's a familiar one: every table full, a fifteen-minute wait list at the door, servers moving fast, and yet the tickets in the kitchen are backing up on the peel rail faster than the oven can clear them. Customers wait forty minutes for a pizza that should take ten. Some walk. Some order fewer extras because they're annoyed rather than relaxed. The owner looks at a packed room and a good top-line number and assumes the business is thriving, when the truthful read is that the site left money on the table by not being able to serve the demand that was already standing in the door. This is the single most common capacity misdiagnosis I see in pizza specifically, because pizza's bottleneck is so often invisible — it's not a missing till, a missing server, or a missing table. It's a box that can only hold so much heat doing so much work per hour, and that number rarely gets calculated until it's already been a problem for months.

Your oven has a hard number: pizzas per hour. That number is not a soft target — it is the absolute ceiling on your site's revenue for any given hour, no matter how many staff you hire, how many tables you have, or how good your marketing is. A dining room with 80 seats and an oven capable of 60 pizzas an hour is, functionally, a 60-pizza-an-hour restaurant wearing an 80-seat costume.

The calculation is **the Peel-to-Plate Rate**, and it's built from three inputs: bake time per pizza (the minutes it physically sits in the oven), deck or floor capacity (how many pizzas fit in the oven simultaneously), and the load/unload overhead — the seconds a skilled pizzaiolo spends getting a pizza in and pulling one out, which matters more than owners think because it's dead time the oven isn't cooking during. A wood-fired Neapolitan oven might bake a pizza in 90 seconds but only fit one, maybe two, at a time, and demand a highly skilled peel operator to keep the load/unload overhead low. A deck oven might bake in 6–8 minutes but fit six pizzas across two decks, and tolerate a less experienced operator. These produce very different throughput numbers that owners rarely calculate explicitly — they just feel the oven "getting slow" on a busy night without knowing the actual ceiling they're hitting.

Illustratively: a two-deck gas oven, four pizzas per deck, 7-minute bake time, with 20 seconds average load/unload overhead per pizza handled by one experienced pizzaiolo working both decks. Eight pizzas cooking simultaneously, each taking 7 minutes plus the overhead time to get it in — in practice this oven, run well, produces around 55–60 pizzas per hour at

full tilt, but only if the pizzaiolo never has to stop to also prep, answer the phone, or fix a mis-fired ticket. Add a second person purely dedicated to loading and pulling and you can push that toward 65–70, because the bottleneck was never really the oven's physics, it was the single human managing it.

This is the insight that should change how you think about capacity spend: a second oven is very often cheaper than more seats. Adding 20 seats to a dining room means more furniture, more floor space (rent), more front-of-house labour, and does nothing for your throughput ceiling if the kitchen still can't produce more pizzas per hour to fill those seats. A second deck oven, by contrast, might cost what three months of the extra rent for that expanded dining room would cost, and it directly lifts your revenue ceiling for every remaining hour you operate the business. I've watched owners spend on the wrong side of this trade-off more than once — a beautiful expanded dining room that still bottlenecks at the exact same 60 pizzas an hour it did before the renovation, because nobody touched the oven.

It's also worth separating the two different capacity problems a growing owner faces, because they call for different fixes. The first is steady-state throughput — can your oven sustain a given pizzas-per-hour rate for three or four hours straight without the pizzaiolo burning out or quality slipping as fatigue sets in. The second is peak-minute burst capacity — the fifteen-minute window at 7:45pm where every table's order lands in the kitchen within moments of each other, which is a queuing problem as much as a throughput one, and which a second, smaller oven dedicated purely to absorbing peak bursts can solve more

cheaply than doubling your main oven's footprint. Several multi-site owners I've watched solve their worst Friday-night backups not by replacing their main oven, but by adding a compact secondary deck oven sized just to shave the top off their fifteen-minute burst, running cold the rest of the week.

Build **the Peak-Hour Capacity Planner** for your own site: take your Peel-to-Plate Rate (pizzas/hour at full effort) and compare it against your actual peak-hour order volume from your POS data over the last eight weekends. If your peak hour is regularly hitting 90%+ of your calculated throughput ceiling, you are not understaffed at the till or the door — you are oven-constrained, and no amount of front-of-house hiring will fix a Friday night that's actually capped by how fast dough can leave a deck. If your peak hour sits comfortably under 70% of ceiling, your bottleneck is elsewhere — likely front-of-house speed, delivery driver availability, or demand itself — and a second oven would be an expensive answer to a question nobody asked.

This week: calculate your own Peel-to-Plate Rate using your actual bake time, deck capacity, and honest load/unload overhead, then pull your last eight weekends of peak-hour order counts from your POS and see how close you're running to the ceiling. If you're above 85%, start pricing a second oven before you start pricing a bigger dining room.

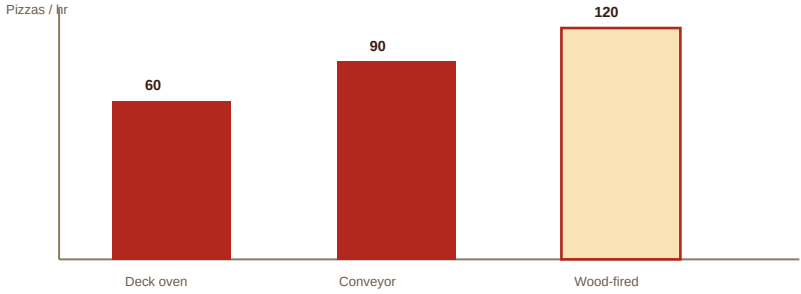

PEEL-TO-PLATE RATE BY OVEN TYPE — THE NUMBER THAT SETS YOUR SITE'S REAL CAPACITY CEILING

CHAPTER 5

Menu Economics Of A Pizza Menu

Pizza menus have a specific way of dying by a thousand small additions. A regular asks for extra mushrooms, you say yes because saying no over a topping feels petty. A new pizzaiolo wants to add their signature creation. A trend article says truffle oil is having a moment. None of these individual decisions feels like it costs anything, and each one, alone, doesn't. Collectively, over two or three years, they turn a tight, profitable, easy-to-execute 14-pizza menu into a sprawling 30-item menu where nobody — not you, not your managers, not your new hires at site three — actually knows which items make money.

I ask owners in this position a simple diagnostic question: if I handed your printed menu to your newest pizzaiolo right now and asked them to name the correct topping weights for every item without looking anything up, how many would they get right? At a single well-run site the answer is usually "most of them," because repetition and proximity to the owner fill the gaps written specs leave open. At a three-site business where that same menu has been copied to two other locations run by people who never watched the original owner build the recipes by feel, the honest answer drops fast — and it drops fastest on exactly the low-volume items that are ordered rarely enough that nobody gets enough repetition to build muscle memory around them, which is precisely the dead-weight quadrant this chapter is about to identify.

The core economics of a pizza menu run through one ingredient more than any other: cheese. Mozzarella is typically the single largest variable cost on your pizza P&L, and it's also the ingredient most vulnerable to what I call **topping creep** — the slow, invisible growth in the average grams of cheese and toppings going onto each pizza as staff turn over and "a generous handful" gets reinterpreted generously each time a new hire watches an old hire do it slightly heavier than spec. A pizza specified at 120g of mozzarella that drifts to 145g over eighteen months, multiplied across every pizza sold, is a margin leak nobody notices on any single ticket and everybody notices on the annual food cost percentage.

The tool is **the Slice Margin Matrix** — every pizza on your menu, plotted on two axes: contribution margin in euros, and sales volume as a share of total pizza sales. This produces four quadrants, and each one tells you something different to do. High margin, high volume: your stars — protect these, never let their spec drift, feature them. High margin, low volume: hidden winners — worth a small push in menu placement or staff recommendation to grow volume. Low margin, high volume: your problem children — these are eating your food cost percentage precisely because customers love them, and they need either a small, careful price adjustment or a spec tightening, not removal. Low margin, low volume: dead weight — the pizzas that exist because someone liked them once, that add prep complexity, inventory SKUs, and menu confusion without earning their place.

Illustratively: a Margherita priced at €11, costing €3.20 in ingredients including 130g of mozzarella at €9.40/kg, delivers a

contribution margin of €7.80 — and if it's your top seller, that's a star. A "Four Meats Deluxe" priced at €16, loaded with 180g of mozzarella plus salami, ham, bacon, and sausage, might cost €7.90 in ingredients — a contribution margin of just €8.10, barely more euros than the Margherita despite a 45% higher price, and if it's also a high-volume item, it's quietly the worst pizza on your menu to sell more of, not the best, purely because the topping cost scales faster than the price does.

This is why a shorter menu consistently outperforms a longer one across multiple sites, and it isn't really an aesthetic argument — it's an execution argument. A 14-item menu means every pizzaiolo at every site can hold every spec in their head accurately. A 30-item menu means specs live on cards nobody checks during a Friday rush, cheese portions drift differently at each site because each site's staff developed their own shorthand for the items they make less often, and your Slice Margin Matrix becomes impossible to keep current because a third of the items sell too rarely to generate a reliable margin picture at all.

Run the Slice Margin Matrix on your full menu, illustratively across, say, 22 items: you will very often find that 6 to 8 of them sit in the dead-weight quadrant, contributing under 3% of total pizza sales combined while carrying a full share of your inventory complexity, prep training time, and spec-drift risk. Cutting those items rarely costs you meaningful revenue — most of that volume migrates to your stars and hidden winners, which are already your best-margin items — while it meaningfully simplifies training at every site you're trying to open.

Sides and drinks deserve a place in this analysis too, because they're frequently the quiet profit engine that subsidises a pizza menu priced for competitiveness rather than pure margin. A garlic knot basket or a side salad often carries a contribution margin percentage well above any pizza on the menu, simply because the ingredient cost is low and customers don't price-compare a side against a competitor the way they mentally price-compare a Margherita. A canned or fountain drink, similarly, routinely delivers a contribution margin north of 70%. None of that shows up if your margin analysis only looks at pizzas — and it means a manager focused purely on pizza sales volume is optimising the wrong number, while a manager trained to suggest a side and a drink with every order is quietly protecting your overall site margin more than any pizza-specific decision could.

This week: build the four-quadrant matrix for your actual menu using your last 90 days of item-level sales data and your true ingredient cost per pizza — not your assumed cost, your actual current cheese and topping weights, weighed if you haven't checked in a while. Identify your dead-weight quadrant and decide, for each item in it, whether it earns its place on a menu you're about to ask a second and third site to execute perfectly.

SLICE MARGIN MATRIX — EXAMPLE MENU

Pizza	Food cost	Menu price	Margin
Margherita	€1.40	€9.50	85%
Pepperoni	€2.10	€11.00	81%
Four cheese	€2.60	€12.50	79%
Prosciutto & rocket	€3.30	€14.00	76%
Menu average target	**—**	**—**	**80%+**

CHAPTER 6

The Platform Problem: Delivery Economics

The first time an owner shows me their delivery platform statement next to their in-house order data, there's usually a long pause. Not because the commission rate is a surprise — most owners can quote their platform's cut from memory, somewhere between 20% and 35% depending on the deal and the market. The pause comes when they see what that commission actually does to a pizza's margin once it's stacked on top of ingredient cost, packaging, and the platform's payment processing fee, and they realise a pizza that looks profitable on the menu can be nearly break-even, or worse, the moment it leaves through a platform.

Here's the arithmetic most owners haven't run explicitly. Take a pizza priced at €14 in-store. Ingredient cost, illustratively, €4.10. Packaging for delivery, another €0.35 you don't pay for a dine-in order. A platform commission of 28% on the €14 comes to €3.92. Add the platform's payment processing cut, often bundled but running roughly 2–3% on top, call it €0.35. Total cost to deliver that pizza through the platform: €8.72, leaving a contribution margin of €5.28 — versus €9.55 for the same pizza sold in-store. The platform pizza earns you 55% of what the in-store pizza earns, on the same ingredients, same labour to make it, same oven time. Run enough volume through that channel without understanding the gap and you can hit a busy, fully booked night

and still post a disappointing week, because "busy" and "profitable" have quietly become different questions depending on which channel the order came through.

The tool is **the Channel P&L Per Order** — the same pizza, the same ingredients, run through every channel your site actually uses (dine-in, pickup, own delivery, and each third-party platform separately, because their commission structures rarely match), so you can see contribution margin by channel rather than a single blended average that hides the problem. Blended averages are dangerous precisely because they can look healthy — a site running 40% of orders through a platform at 55% of full margin, blended with 60% of orders at full margin in-store and pickup, still shows an overall margin that looks acceptable on the monthly P&L while masking that nearly half your growth in order volume is coming through the channel that pays you the least per pizza.

This doesn't mean quit the platforms — that would be the wrong lesson, and I want to be clear about that, because reach is real and new customer acquisition through a platform is often the only affordable way a new second or third site gets discovered at all in its first few months. What it means is running platforms deliberately, for the job they're actually good at, rather than treating them as your primary channel by default because that's what happened without anyone deciding it. Platforms are excellent at reach — a customer in a new neighbourhood discovering your second site exists, browsing a category page, trying you for the first time with low commitment. They are expensive as a repeat-customer channel, because every single

order from a loyal regular still pays that same 28% commission, forever, unless you do something about it.

That "something" is **the direct-ordering migration plan**: a deliberate sequence for moving a customer from their first platform order to a direct relationship with your site, without pretending the platform doesn't exist. Step one, every box that goes out through a platform still carries a card or a printed insert offering a modest, genuine incentive — say 10% off — for ordering direct next time, because the platform's terms typically prohibit discounting on their own channel but rarely prohibit you promoting your own. Step two, your own ordering system (website or app) needs to be at least as fast and as easy as the platform's, because a customer who tries direct ordering once and finds it clunky reverts to the platform permanently and tells you nothing about why. Step three, track it — the share of repeat customers who've made the switch from platform-first to direct-first ordering is a number worth watching monthly per site, because it's the single best predictor of whether your margin structure is trending toward the platform's terms or your own over time.

Illustratively, a site launching this migration deliberately — insert cards, a genuinely fast direct ordering flow, and staff trained to mention it at pickup — might move from 65% platform-order share to 48% platform-order share over eight months, without any drop in total order volume, because the orders didn't disappear, they just moved to a channel paying nearly double the margin per pizza. On a site doing 1,800 pizzas a month, shifting even 15 percentage points of volume from platform to direct, at

the roughly €4.30 margin gap calculated above, is worth in the region of €1,160 a month in contribution margin recovered — money that was always there, just being paid out the door to a platform commission instead of staying in the business.

One more distinction worth making explicit, because owners conflate it constantly: not all platform commission structures are equal, and multi-site owners have leverage single-site owners don't. Once you're operating three, four, five sites through the same platform account, you have a genuine negotiating position — real, comparable volume data across multiple locations that a platform account manager can use internally to justify a lower blended rate, especially if you're willing to commit to featured placement or exclusivity on one platform over a competitor. Few independent owners ever ask. I've watched multi-site pizzerias secure a two-to-four percentage point commission reduction simply by presenting their own Channel P&L Per Order data to a platform rep and asking directly, armed with numbers rather than a vague complaint about fees being high — the difference between "your commission feels expensive" and "here is exactly what it costs me per order across four sites, and here is the volume I can commit to at a better rate" is the difference between being ignored and getting a meeting.

This week: pull one of your top-selling pizzas and run the Channel P&L Per Order across every channel your site actually sells through, using your real commission rates and real packaging cost. If you don't already have a direct-ordering incentive going out in every platform box, write the insert card

copy today — it costs almost nothing and it's the first step of a migration plan that compounds for years.

CHAPTER 7

Own Delivery, Platform Delivery Or Pickup Only

Every fulfilment decision in this business gets framed as a distance question, and that's the wrong frame. Owners ask "how far should we deliver?" as if distance alone decides profitability, and set a delivery radius on a map — three kilometres, five kilometres — based on gut feel about how far is "reasonable." The actual driver of delivery profitability isn't distance. It's density: how many deliverable orders sit inside whatever radius you draw, because a delivery driver's cost is dominated by time, and time is dominated by how far apart consecutive drops are from each other, not by how far any single drop is from the store.

This is also the decision owners most often make once, at the first site's opening, and then never revisit — even as the site's order pattern changes completely over two or three years. A radius set on opening-week guesswork gets copied verbatim into every subsequent site's setup, including onto platform listings, where a radius that's too wide means drivers spending most of their paid hour driving rather than delivering, and a radius that's too narrow means leaving genuinely profitable, dense nearby demand unserved because nobody thought to widen it once the neighbourhood's ordering pattern became clear from real data.

Picture two five-kilometre radii. One sits over a dense residential area with narrow streets and apartment blocks — a driver can

realistically complete four drops an hour inside it. The other stretches out into lower-density suburban or semi-rural housing with long driveways and wide-spaced streets — the same driver, same vehicle, same radius size, might manage two drops an hour. Identical radius, identical fuel cost per kilometre roughly, wildly different economics, because the second scenario pays a driver's full hourly wage for half the delivery throughput.

The tool is **the Density-Radius Model**: instead of drawing a radius first and hoping for density, you calculate the minimum order density — orders per square kilometre per peak hour — your delivery model needs to break even on driver cost, and then check whether your actual or projected order pattern clears that bar before you commit to own-delivery at all. Start from your fully loaded driver cost per hour, including wages, vehicle cost or mileage allowance, and insurance — illustratively, €19/hour all-in for an employed driver with a company scooter. If that driver can complete, at your site's real street density and speed limits, 3.5 drops an hour, your cost per delivery is €5.43 before a single euro of platform commission is avoided. Compare that to your platform's delivery fee to you as a merchant, or the commission uplift platforms charge for their own delivery — if platform delivery is costing you the equivalent of €4.80 per order and your own-delivery model, at your actual density, costs €5.43, running your own fleet isn't automatically the win it feels like emotionally; it's a decision that needs the real density number, not the assumption that "our own driver must be cheaper than the app."

This is precisely why the same delivery decision can be correct at one site and wrong at another for the same brand. A first site

in a dense urban core might run profitable own-delivery at a 2.5km radius with four drops an hour easily achievable. A second site in a lower-density suburban location, opened because the rent was attractive, might only support 2 drops an hour across the same radius — meaning own-delivery costs nearly double per order at that site, and the right call there is pickup-forward with platform delivery covering the rest, even though it's the "same brand" making "the same pizza."

Vehicle choice matters more than owners give it credit for, too, and it interacts directly with density. A scooter or bicycle courier can weave through the exact narrow streets and awkward parking that make a car slow and expensive in a dense urban catchment, often achieving a higher drops-per-hour rate than a car would in the same footprint despite carrying fewer orders per run — while a car or small van earns its keep in a lower-density, longer-haul catchment where a single multi-order run covering more ground beats several short scooter trips. Match the vehicle to the density profile you calculated, not to whichever vehicle happens to be sitting in the car park.

Packaging density matters too, and it's a lever owners underuse. A driver who can carry three orders in one insulated bag run, because they happen to route through the same building or adjacent streets, effectively doubles their per-hour throughput compared to three separate single-order trips. Route-batching software that groups nearby orders into a single run — rather than dispatching strictly first-in-first-out — is one of the highest-leverage, lowest-cost changes an owner running their own

delivery fleet can make, because it directly attacks the density problem without needing a single new customer.

Illustratively, run the Density-Radius Model for a hypothetical second site: peak-hour order volume of 22 orders across a 4km² catchment, giving a density of 5.5 orders/km²/hour. At your driver's achievable speed and stop time, that density supports roughly 3.2 drops per driver per hour — comfortably above your €19/hour driver's break-even of roughly 2.9 drops/hour needed to beat your platform's delivery fee. Own delivery clears the bar here. Run the same model on a site with 9 orders across the same 4km² catchment — density of 2.25 orders/km²/hour, well under the threshold — and pickup-plus-platform is very likely the more profitable call, even though it feels like a step backward from "having our own drivers" as a brand statement.

There's a hybrid option worth naming explicitly, because owners often treat this as a binary choice when it isn't one: run own-delivery inside a tight, high-density inner radius where the drops-per-hour math clearly wins, and hand off everything beyond that ring to a platform's delivery network, which absorbs the low-density, long-haul orders your own driver would lose money on anyway. This gets you the margin benefit of own-delivery where it actually pays, and the platform's larger driver pool for the orders where distance alone makes any single-driver economics weak, without forcing an all-or-nothing decision across your entire catchment. Several multi-site owners I've worked with run exactly this split, redrawing the inner ring every six months as their own order density data updates.

This week: map your actual peak-hour delivery orders against your radius for each site and calculate real orders per square kilometre per hour, then compare the drops-per-hour your driver can realistically achieve against the driver cost needed to beat your platform's delivery economics. Don't set a radius by feel again — set it by the density number.

CHAPTER 8

The Box Is The Product In Delivery

A pizza that leaves your oven perfect and arrives at a customer's door thirty minutes later disappointing is not a cooking failure. It's a packaging failure, and it's one of the most under-invested problems in the entire delivery chain, because the box gets chosen once, early, usually on price, and then never revisited even as the rest of the delivery operation matures around it.

This is also the review-quietest problem in your entire operation, and that's exactly why it survives unaddressed for years. A customer whose pizza arrived soggy rarely writes "the box venting was inadequate for a 62% hydration dough at a 25-minute travel time." They write "pizza was cold and soggy, wouldn't order again," and most owners read that as a speed problem or a driver problem, and respond by trying to shave minutes off delivery time — which helps, but treats a symptom while leaving the actual packaging defect untouched. I've seen owners spend real money on a second driver to cut average delivery time by four minutes, when a €0.09 packaging change would have solved the same complaint more completely and far more cheaply.

Steam is the enemy, and it's a predictable one. A hot pizza generates moisture as it sits in a sealed box — the crust's residual heat drives water vapour off the sauce and cheese, and if that vapour has nowhere to go, it condenses on the inside of the lid and drips back down onto the pizza, turning a crisp base

soft within minutes. This is why a pizza that was genuinely excellent when it left your oven can arrive fifteen minutes later with a bottom crust that has gone from crisp to steamed-limp, through no fault of the recipe, the dough, or the driver — purely a function of box design and travel time interacting badly.

The fix lives in venting, and venting is a design problem with a right answer that depends on your specific box, your specific pizza's moisture content, and your specific typical travel time — which is exactly why generic advice like "use a vented box" undersells the work. Some boxes need cut vents on the lid. Some benefit from a corrugated liner that lifts the pizza slightly off the box floor, letting air circulate underneath rather than trapping moisture directly against the crust. Some benefit from both. None of this is guessable from a supplier's marketing copy — it needs to be tested against your actual product and your actual delivery times.

That's **the Fifteen-Minute Box Test** — a simple, repeatable protocol every site should run whenever it changes box supplier, box design, or meaningfully changes its average delivery time. Bake a standard pizza to your normal spec. Box it exactly as you would for a real delivery order. Let it sit, sealed, for the duration of your site's actual average delivery time — not your best-case time, your real average, pulled from your dispatch data. At the end of that time, open the box and score it against four criteria: crust texture (crisp, acceptably firm, or soft/soggy), visible condensation on the lid, cheese and topping slide (has anything shifted position, a sign of insufficient venting or an underbaked base that didn't hold structure), and temperature at the centre of

the pizza with a probe thermometer. Run this test with three or four candidate box-and-vent configurations side by side, scored by more than one person so you're not just confirming your own bias toward whatever box you already bought a pallet of.

Illustratively: a site testing three configurations at its real 22-minute average delivery time found its existing unvented box scored soggy on crust texture with heavy lid condensation; a two-vent-hole version of the same box scored acceptably firm with light condensation; and a vented box with a corrugated liner scored crisp with no visible condensation at all. The liner version cost €0.09 more per box than the plain vented version — on 1,800 delivery orders a month, that's an extra €162 monthly spend. Against that cost, the site's delivery-specific review complaints mentioning "soggy," "cold," or "fell apart" — tracked as a simple keyword tag on incoming reviews and refund requests — dropped by roughly two-thirds over the following quarter. At even a conservative estimate of what one avoided one-star review costs a site in lost future bookings, €162 a month is a trivially easy trade.

Travel time interacts with all of this more than owners expect, which loops straight back to the density model from the last chapter — a site running tight, dense, four-drop-an-hour delivery routes is protecting its box's performance almost as much as it's protecting its driver economics, because every extra minute in the box is another minute of steam working against the crust. A box that scores "acceptably firm" at your average 18-minute delivery time might score "soggy" at 32 minutes if your routing is inefficient on a busy night, which means your packaging choice

and your delivery radius decision from Chapter 7 are not actually separate decisions — they're the same decision viewed from two different angles, and testing your box only at best-case delivery times will hide a problem that shows up precisely on your busiest, most profitable nights.

Multi-site owners have one more reason to run this test regularly that single-site owners don't: box performance is not guaranteed to travel across sites even when the box itself is identical, because ambient kitchen humidity, oven type, and average delivery time can all differ site to site. A commissary-supplied dough (Chapter 3) baked in a different oven type at a second site can hold moisture differently than the same dough baked in your original oven, which means a box configuration validated at site one is a starting hypothesis at site two, not a guaranteed answer — run the Fifteen-Minute Box Test independently at every site rather than assuming the first site's result transfers.

This week: run the Fifteen-Minute Box Test — or better, the test at your actual average delivery time, whatever that number is — against your current box and at least one alternative configuration, scored by two people independently. If nobody at your site has done this in the last twelve months, your box has almost certainly never been validated against how your delivery times have actually changed as you've grown.

CHAPTER 9

Site Selection For Pizza

The most expensive mistake in this whole book is picking the wrong second location, because unlike a bad dough batch or an underperforming menu item, a bad site is a multi-year lease commitment you can't fix with better training or a tighter spec. And the mistake almost always comes from the same root cause: the owner evaluates the second site as real estate — foot traffic, rent per square metre, how the space looks — rather than as a pizza production and delivery facility with its own specific requirements that have nothing to do with how charming the storefront looks.

Four factors matter more for a pizza site than for almost any other restaurant concept, and each one gets under-checked in the excitement of finding a space that "feels right." Residential density within your realistic delivery radius, because pizza delivery economics live and die on the density math from Chapter 7, and a beautiful high-street location surrounded by offices that empty out at 6pm is a poor delivery site no matter how good the foot traffic looks during lunch. Footfall for walk-up and slice trade, if that's part of your model, which is a genuinely different metric from residential density and needs its own check — a site can have superb walk-by counts and terrible surrounding household density, or vice versa. Direct competition, meaning not just "is there another pizzeria on this street" but a more specific question: is there another pizzeria of comparable

quality and price point already serving this exact catchment, because a site with three mediocre chain pizza options nearby is a much better opportunity than a site two doors down from another genuinely excellent independent. And gas or power supply sufficient for your oven — a factor so basic it gets skipped, and so expensive to fix after signing that I've watched owners spend five-figure sums on an electrical or gas main upgrade they didn't know they'd need until the oven installer showed up.

The tool is **the Pizza Site Score** — a weighted checklist run before you sign anything, not after. Score each candidate site 0–5 on: residential density within delivery radius, walk-up footfall (if relevant to your model), direct competition intensity (5 = little to none, 0 = saturated with comparable quality), utility capacity for your oven type confirmed in writing by the landlord or a contractor visit rather than assumed, parking or loading access for deliveries and driver pickup, rent as a percentage of your realistic projected revenue for that specific catchment rather than a citywide average, and — the factor owners skip most often — visibility and signage rights, because a site tucked behind another building with restricted signage can have great underlying density and still underperform for years on discovery alone. Weight utility capacity and residential density most heavily, since those two are the hardest and most expensive to fix after the fact; weight rent-to-revenue and competition next; treat footfall and signage as real but secondary factors that matter more for a slice-forward format than a delivery-forward one.

Here's the principle that should guide every second-site search, and it's the one most owners resist because it feels unambitious: the second site should be a copy of the first, not an upgrade. The instinct after a successful first location is to reach for something bigger, flashier, in a more prestigious part of town, because success feels like it should be rewarded with a nicer space. Resist it. Your first site succeeded because of a specific combination of density, competition, and utility fit that you can now describe precisely using the Pizza Site Score — go find another location that scores similarly on that same profile, in a different catchment, rather than a location that scores differently but "feels like an upgrade." A bigger, more prestigious space with worse residential density and a tighter power supply is not a better site. It's a worse site with nicer flooring.

Illustratively: a first site scores 4 on residential density, 3 on footfall, 4 on competition, 5 on utility capacity (purpose-fitted from the start), 3 on parking, 4 on rent-to-revenue, 3 on visibility — a total of 26 out of 35. A candidate second site the owner is excited about, in a trendier part of town, scores 2 on density (mostly offices), 5 on footfall, 2 on competition (two other well-regarded independents nearby), 2 on utility (would need a gas main upgrade), 4 on parking, 2 on rent-to-revenue (rent is 40% higher than the first site relative to realistic projected volume), and 5 on visibility — a total of 22 out of 35, and critically, weak on exactly the two heaviest-weighted factors. A second candidate in a quieter, more residential neighbourhood scores 5, 2, 4, 4, 4, 5, 2 — a total of 26, matching the first site's profile almost exactly despite feeling, on a walk-through, far less exciting than the

trendy option. The Pizza Site Score exists precisely to catch this gap between what feels exciting and what will actually perform.

One practical note on the utility factor specifically, because it's the one most likely to be assumed rather than confirmed: never take a landlord's word for gas or electrical capacity, and never assume a space that previously housed a restaurant automatically has what your oven needs. A previous tenant's fryer-and-grill setup tells you nothing about whether the electrical supply can handle a deck oven's draw, and a previous tenant's residential-grade gas connection may be entirely inadequate for a commercial pizza oven's BTU requirement. Get a contractor to physically confirm capacity against your specific oven's spec sheet before you sign, not after — the cost of that visit is trivial next to the cost of discovering the gap after your lease is signed and your equipment is on order.

This week: score your existing best-performing site on the Pizza Site Score to establish your own baseline profile, then score every site currently under consideration against that same baseline — and be suspicious of any candidate that excites you but scores meaningfully lower on density or utility capacity than the site you're trying to replicate.

CHAPTER 10

The Manager Who Is Not You

This is the chapter most owners resist reading carefully, because it's not really about pizza. It's about the hardest psychological transition in growing any owner-operated business: learning to trust someone else with the thing you built with your own hands, at a standard you may never fully believe matches your own, because it probably won't — and the business has to succeed anyway.

I've sat with owners who can recite their dough hydration percentage to the decimal point but freeze completely when asked what, specifically, they expect their new site manager to decide alone versus escalate to them. That vagueness isn't a personality flaw — it's the natural result of never having needed to articulate it before. At one site, the owner was the decision-maker for everything, so there was never a boundary to define. At two sites, the absence of that boundary becomes the single biggest source of friction between owner and manager, because the manager either escalates everything (exhausting the owner and defeating the purpose of hiring a manager at all) or decides everything alone (terrifying the owner, who finds out about problems only after they've compounded).

The fix is **the Manager Mandate** — a written document, given to every site manager on day one, that draws the line explicitly rather than leaving it to be discovered through trial and error over painful months. It has three lists, not one. First, standardised and

non-negotiable: the items from your Dough DNA Sheet, your Slice Margin Matrix specs, your opening and closing checklists, your health and safety procedures — the things that must be identical at every site regardless of manager preference, because these are what protect the brand and the Copyability Scorecard investment from Chapter 1. Second, delegated and manager-owned: staff scheduling within budget, day-to-day supplier ordering within agreed par levels, handling customer complaints up to a defined refund threshold, local marketing decisions within a monthly budget — decisions the manager should make without checking in, because requiring approval on these erodes both their authority and your own time. Third, escalate-immediately: anything involving a health and safety incident, a significant negative review or public complaint, a request from a supplier to change terms, staff disciplinary action beyond a verbal warning, or any spend above an agreed threshold — the handful of situations where the cost of a delay is lower than the cost of a manager making the wrong call alone.

Illustratively, a Manager Mandate might set the refund threshold at €40 — below that, the manager resolves it on the spot with an apology and a comp, no owner call needed; above it, or for anything involving a public review, the owner is looped in within 24 hours. It might set the marketing budget at €150/month, spendable by the manager on local flyers, a school fundraiser sponsorship, or a social media boost, without needing sign-off for each individual spend. Numbers like these feel small to write down and enormous in effect, because they convert dozens of weekly "can I just check with you" moments into decisions the

manager makes confidently and immediately, which is exactly what a customer standing at the counter needs in the moment a problem happens.

The second half of this chapter's tool is **the weekly review format** — a fixed-agenda, time-boxed call or meeting between owner and each site manager, ideally the same day and time every week so it becomes a rhythm rather than an ad hoc "checking in" that happens when the owner remembers or worries. Keep it to four items, in this order: the Multi-Site Scorecard numbers from Chapter 14 for the week just closed, compared against target and against the site's own trailing average; any escalate-immediately items from the Manager Mandate that occurred; one thing that went well worth naming specifically, because a review that's only ever about problems trains managers to hide small ones before they become big ones; and one specific action for the coming week, agreed by both people, not assigned unilaterally by the owner. Thirty minutes, same structure every time, and critically — not a call where the owner walks through every ticket from the week asking why things happened, because that behaviour teaches a manager to stop bringing problems to you voluntarily.

I want to be direct about the emotional cost of this chapter, because pretending it's purely mechanical would be dishonest. Handing over the peel, the ordering decisions, the way a customer complaint gets handled — all of it — to someone who is not you is going to produce moments where you watch a choice get made that you wouldn't have made yourself, and it will not be a catastrophe, and it will still feel like one the first several

times. The Manager Mandate doesn't remove that feeling. It gives you a structure to check whether the outcome that worried you was actually a mandate violation (which needs a direct conversation) or just a legitimate decision inside the manager's delegated zone that simply wasn't the decision you'd have made (which needs you to let it go), and learning to tell those two things apart quickly is most of what makes an owner capable of running five sites instead of one.

One warning sign to watch for as you hand over control: a manager who never escalates anything is not necessarily a great manager operating confidently inside their mandate — sometimes they are a manager quietly absorbing problems they're afraid to bring to you, which shows up months later as a resignation you didn't see coming, or a standard that's drifted further than you realised because nothing ever surfaced. The weekly review's "one thing that went well" question exists partly to build the trust that makes the harder disclosures possible — a manager who has never once heard genuine, specific credit from you is unlikely to be the one who calls you at 9pm to say a health inspector just walked in unannounced.

This week: write the three-list Manager Mandate for your business — standardised, delegated, escalate-immediately — with real numbers on every threshold, and schedule the first fixed-format weekly review with each site manager you currently have.

CHAPTER 11

Training Pizzaiolos At Scale

Every pizzeria owner has, at some point, depended entirely on one person. Usually it's the first pizzaiolo hired, or the owner themselves — someone whose hands know the dough, whose eye can judge a bake by colour alone, whose speed on the peel sets the pace for the whole kitchen. That dependence feels like an asset right up until the day that person calls in sick on your busiest Saturday, or gives notice to open their own place, and the whole operation's quality visibly wobbles because the skill was never actually in the business — it was in one person who happened to work there.

This is the single-star-pizzaiolo trap, and it gets worse, not better, as you add sites, because now you need that same irreplaceable skill in three or four places simultaneously, and there is exactly one of the person who has it. The instinct to solve this by hiring "another star" — poaching an experienced pizzaiolo from a competitor for each new site — is expensive, slow, and doesn't actually fix the underlying problem, because you've just created three or four single points of failure instead of one.

The real fix is building a skill progression that turns pizza-making from an art one person holds into a craft multiple people can be trained into, systematically, with clear checkpoints. That's **the Pizzaiolo Ladder** — a five-stage skills matrix that breaks the full skill set into stages a new hire can climb in a defined order, each

with an explicit time-to-competency target and a pass/fail standard rather than a vague sense of "getting better."

Stage one is dough handling: portioning, balling, and basic stretch technique, without touching a hot oven. A motivated new hire should reach a consistent, acceptable standard here within one to two weeks of supervised practice, judged against a simple checklist — even ball weight within a 5g tolerance, no tears in the stretch, consistent thickness across the disc measured at four points. Stage two is topping and dressing: correct portioning against your Slice Margin Matrix specs from Chapter 5, correct placement, correct order of operations (sauce before cheese, cheese before certain toppings, some toppings added post-bake) — typically another one to two weeks, and the stage where topping creep either gets prevented or gets baked in as a bad habit for years. Stage three is the peel: loading a stretched, topped pizza onto the peel without deformity, launching it cleanly into the oven, and — this is the part that takes longest — reading the bake in progress and knowing when and how to rotate it. Three to five weeks typically, and the first stage where a genuine skill gap between trainees becomes visible, because peel work rewards repetition and hand-eye coordination in a way the earlier stages don't as strongly. Stage four is pace: running peel and oven simultaneously during real, busy service, holding quality standard under time pressure rather than only in a quiet training kitchen — this is where a technically competent trainee often reveals they can't yet hold the standard once four tickets land at once, and it typically needs four to eight weeks of real service exposure to solidify. Stage five is training others: the final stage,

and arguably the most valuable one for a multi-site business specifically, because a pizzaiolo who can teach stage one through four to a new hire is the person who lets you open a fifth site without needing to poach anyone.

Each stage ends in **the certification checklist** — a specific, observable pass/fail test, not a manager's subjective impression. Illustratively, the stage-three peel certification might require a trainee to load and launch ten consecutive pizzas with zero deformities, correctly identify and rotate at least eight of ten pizzas at the right moment based on visible bake progress, and complete the full ten-pizza sequence within a defined time window that matches your site's real service pace. Pass all three and they're certified at stage three; fail any one and they get another week of supervised practice before retesting, rather than being quietly waved through because the shift was busy and the manager didn't want to deal with a difficult conversation.

The certification record matters for a second, less obvious reason beyond training quality: it's your actual staffing capacity map across sites. When you know precisely how many stage-three-and-above certified pizzaiolos you have across your whole business, and where they're currently working, you can see staffing risk before it becomes a Saturday-night crisis — a site with only one stage-four-certified person on the roster is one sick day away from a quality collapse, and that's a fact you want to know on a Tuesday planning meeting, not learn the hard way on a Friday at 7pm.

Illustratively, a three-site business running the Pizzaiolo Ladder deliberately might track: Site A, four stage-four-certified, two stage-five (training-capable); Site B, two stage-four, zero stage-five; Site C, one stage-four, zero stage-five. That map tells you immediately that Site C is fragile — a single point of failure exactly like the one-star-pizzaiolo trap this chapter opened with, just wearing a job title instead of a name — and that your two stage-five trainers at Site A are an underused asset who should be spending time at Site C, not permanently parked where they already have backup.

The economics support the time investment clearly. A business relying on poaching experienced hires to staff new sites typically pays a wage premium of 15–25% above what a home-grown, ladder-trained pizzaiolo costs, plus a slower ramp because a poached hire still needs to unlearn a previous employer's habits and relearn your specific Dough DNA Sheet and Slice Margin Matrix specs — often taking nearly as long to reach full site-specific competency as a motivated new hire trained from stage one on your own system from day one.

This week: map your current pizzaiolo team against the five-stage ladder honestly, by site, and identify any site with fewer than two stage-four-or-above certified people — that's your immediate single-point-of-failure risk, and it's worth fixing before you open a fourth or fifth site rather than after.

CHAPTER 12

Purchasing Across Sites

A single-site owner's purchasing relationship with a flour mill, a cheese supplier, or a box manufacturer is usually simple and personal — a standing weekly order, a familiar rep, a price that hasn't been seriously renegotiated in years because nobody had leverage to renegotiate it with. Add a second and third site and something changes that most owners don't immediately notice: you now have real purchasing volume, and real purchasing volume is worth something to a supplier that a single site's order never was.

The mistake I see most often at this stage isn't failing to consolidate purchasing — it's consolidating too narrowly, around price alone, without building in the quality control and supplier-risk thinking that should come with it. An owner discovers that ordering flour centrally for all three sites gets a meaningfully better price than each site ordering separately, switches immediately, and stops there — without noticing that centralising also means every site's dough now depends on one delivery schedule, one supplier relationship, and one product spec, which is a different kind of risk than the one it solved.

The purchasing consolidation plan handles this properly by separating "what to consolidate" from "how to protect against what consolidation creates." Start with the categories where consolidation has the clearest upside and the least added risk: flour, canned tomatoes, and boxes are usually the safest early

wins, because they're commodity-adjacent products where multiple qualified suppliers exist and switching cost if one supplier fails you is genuinely low. Mozzarella and other fresh dairy deserve more caution, and I'll come back to why in a moment. Build the plan category by category, negotiating a volume-tiered price schedule with each supplier — a price that steps down as your combined multi-site order volume crosses defined thresholds — rather than a single flat discount, because a tiered structure gives you a clear, visible incentive to keep consolidating as you add sites four, five, and six, and gives the supplier a clear reason to want to keep earning your growing volume rather than treating your current order as a ceiling.

The quality control benefit of consolidation is real and worth stating plainly, because it's easy to focus purely on price and undersell this part: when every site orders the same flour from the same mill on the same specification sheet, you've removed an entire category of the recipe-drift problem from Chapter 1 before it can even start, because the input itself is now identical everywhere rather than "similar, sourced locally, close enough." This is quietly one of the highest-leverage moves available to a growing pizza business, because it fixes a consistency problem at the source rather than trying to train it out of variable inputs at each site.

Now the risk side, because it's the part consolidation plans skip. **The Single-Source Risk Score** exists to stop you from accidentally building your entire multi-site quality standard on top of one supplier relationship that could fail you with no warning — a mozzarella producer having a bad month, a flour mill having a

fire, a box manufacturer going under. Score each key ingredient or supply category on three dimensions, 0–5 each: switching difficulty (how hard would it be, in practical terms and in time, to qualify a second supplier for this product to your spec if the first one failed tomorrow), lead time risk (how much notice would you realistically get before a supply problem hit your kitchens), and quality sensitivity (how much would the finished pizza suffer if you had to substitute a lower-quality version of this input temporarily while you sourced a replacement). A high total score — meaning hard to switch, little warning, and highly sensitive to substitution — flags a category where single-sourcing for a better price is a genuine strategic risk, not just an efficiency win.

Mozzarella is the category that scores highest on this almost everywhere I've seen it applied, and it's why I flagged it above. It's the ingredient your pizza is most sensitive to in terms of final quality — moisture content, melt behaviour, and flavour vary meaningfully between producers in ways customers genuinely notice, even if they couldn't articulate why one pizza tasted better than another. It's also often produced by relatively few qualified suppliers at the volume and consistency a multi-site pizzeria needs, meaning lead time and switching difficulty both score high too. The sensible move for a category that scores high across all three dimensions of the Single-Source Risk Score isn't to avoid consolidating it — the quality control benefit is too valuable to skip — it's to maintain a qualified, tested second supplier on standby even while your primary volume runs through your main relationship, so a disruption becomes an inconvenient phone call rather than an existential one.

The supplier scorecard completes the system: a simple quarterly review, per key supplier, on price competitiveness against market, on-time delivery percentage, product consistency (measured against your own spec sheet, not the supplier's own claims), and responsiveness when something goes wrong. Illustratively, a mozzarella supplier scoring strong on price and consistency but weak on on-time delivery — say, 91% on-time against a target of 98% — is a supplier worth a direct conversation before the relationship sours into a crisis, and a documented scorecard trend gives you the specific, dated evidence to have that conversation productively rather than emotionally.

Timing your consolidation push matters too. The best moment to renegotiate a supplier relationship is not when you're desperate for a better price after a rough quarter — it's when you're about to open a new site and have a genuine, dated volume increase to offer as leverage. Bring the supplier a specific number: "we're adding a fourth site in Q2, which moves our monthly flour volume from 2.4 tonnes to roughly 3.2 tonnes" is a conversation that gets you a real tiered discount, where "can you do better on price" without a concrete volume story rarely does.

This week: run the Single-Source Risk Score against your three or four highest-volume ingredient categories, and for any category scoring high across all three dimensions, find out today whether you actually have a qualified backup supplier or whether you've been assuming one exists without ever having confirmed it.

CHAPTER 13

One Brand, Three Neighbourhoods

A friend of mine — well, an owner whose business runs on our platform, which is close enough to a friend after enough late-night phone calls — once described opening his third site as "the moment I realised the brand wasn't a recipe, it was a promise, and I'd never actually written down what the promise covered." He'd spent two years assuming everyone, including himself, already knew where the line sat between "this is who we are everywhere" and "this can look different here." Nobody did, and the argument that finally forced the question wasn't a big one — it was a manager who wanted to swap the house salad dressing for a local favourite, a genuinely small decision that turned into a two-week disagreement purely because there was no document either side could point to.

Two mistakes sit on opposite sides of the same question, and I've watched owners fall into each one with equal frequency. The first is the brand-consistency trap: forcing every site to be absolutely identical down to the last detail, regardless of the neighbourhood it sits in, because consistency feels safe and any local variation feels like a slippery slope toward losing control of the brand. The second is the local-irrelevance trap: letting each site drift so far toward its own neighbourhood's preferences that the three locations barely feel like the same business anymore, and the brand you spent years building at site one has quietly

dissolved into three unrelated local restaurants that happen to share a sign.

Both traps come from the same unexamined assumption — that consistency and local relevance are opposites, and you have to pick a point on a single dial between them. They aren't opposites. The actual skill is deciding, deliberately and in writing, exactly which elements of your business must be identical everywhere and which elements are genuinely fine, even good, to flex by neighbourhood — and then holding that line firmly in both directions, rather than letting it drift by accident at either end.

That's **the Core-and-Flex Rule**, documented as a brand rules sheet with two explicit lists, given to every site manager and referenced every time a "should we do this differently here" question comes up. The fixed list covers everything that defines what your pizza actually is and what your brand actually promises: your Dough DNA Sheet specs from Chapter 2, your core menu items and their Slice Margin Matrix specs from Chapter 5, your logo, colour palette and signage standards, your pricing philosophy even if not identical euro figures, your customer service standards, and your food safety and hygiene procedures. None of this changes by neighbourhood, ever, because a customer who's had your Margherita at site one and orders it at site three is trusting that it's the same pizza, and breaking that trust for the sake of local flair costs you more than the flair is worth.

The flex list covers everything that can genuinely improve by adapting to a specific neighbourhood without threatening what the brand promises: a small number of local special items — one or two rotating specials that use local or seasonal ingredients, clearly marked as that site's own addition rather than presented as core menu; local marketing and community involvement, which should look different in a university-area location than in a family-suburban one; interior styling details within your fixed brand palette; hours of operation matched to genuinely different local demand patterns; and delivery-versus-dine-in emphasis matched to the site's actual density profile from Chapter 9, which might mean a smaller dining room and bigger kitchen at one site and the reverse at another.

Illustratively, a three-site pizzeria's brand rules document might fix eleven core menu items with locked specs and pricing philosophy identical across sites, while allowing each site up to two rotating local specials that must still use the same base dough and same core sauce, so even the "local flex" item still tastes unmistakably like the brand underneath its local twist. One site near a university runs a late-night student special two nights a week; another in a family-residential area runs a Sunday family-bundle deal; the third, in a business district, skips both and instead runs a fast weekday lunch special built for a 20-minute table turn. All three still serve the identical Margherita at the identical spec, because that's the thing customers are actually trusting the brand to deliver consistently, wherever they encounter it.

The test I give owners for sorting a new proposal into fixed or flex is simple: would a regular customer from site one feel confused or misled if they discovered this difference at site three? A different rotating special, different hours, a different community sponsorship — no confusion, no problem, genuinely good local relevance. A different dough recipe, a different core pizza's toppings, a materially different price for the same item with no stated reason — that would feel like a bait-and-switch to a customer who trusted the brand across locations, and it belongs on the fixed list even if the site manager has a genuinely good local reason to want to change it.

One practical warning: review the flex list itself periodically, not just compliance with it, because local flex items have a tendency to accumulate complexity over time in exactly the way Chapter 5's topping creep does — a rotating special that started as one clean seasonal item can, eighteen months later, have quietly spawned three variants, none of which anyone remembers approving. Put an expiry date on every local special when it's approved, forcing a deliberate renewal decision rather than allowing indefinite quiet accumulation.

The same document also settles arguments before they start, which is worth more than it sounds. A written Core-and-Flex Rule, agreed and shared before a new manager's first difficult local-adaptation request lands, turns what could become a personal disagreement about taste or judgment into a simple lookup — is this on the fixed list or the flex list — and lets the manager answer their own question most of the time without needing you on the phone at all.

This week: write your own Core-and-Flex Rule as two explicit lists, share it with every site manager, and audit your current locations for anything that's drifted onto the flex side of the line without ever having been deliberately approved to be there.

CHAPTER 14

The Numbers That Compare Sites

The hardest management question in a multi-site pizza business isn't "how is the business doing" — your monthly P&L answers that well enough. It's "which site needs my attention this week, and why," and most owners answer it by feel: whichever site called with a problem most recently, or whichever site the owner happened to visit last and noticed something off. That's not management, it's whack-a-mole, and it means the sites that don't call — often because a manager is quietly struggling and doesn't want to admit it, exactly the risk flagged in Chapter 10 — go unexamined for months at a time.

There's a specific failure pattern I want to name because it's so common it deserves its own label: the loudest site is rarely the site that needs you most. A manager who calls often, escalates small things, and visibly needs reassurance is usually a manager who's engaged enough to notice problems — which is a good sign, even if the constant calls feel exhausting. The genuinely dangerous site is the quiet one, run by a manager who's proud, or new, or conflict-avoidant, absorbing a slow decline without ever raising a hand, until the scorecard — or a customer's public review — surfaces it months after it started. A weekly, numeric, comparable measurement is the only reliable defence against that pattern, because it doesn't depend on anyone deciding to speak up.

The Multi-Site Scorecard fixes this by giving every site the same weekly measurement on the same day, so sites can be compared fairly against each other and against their own trailing average, rather than compared on gut feel or on whoever's been loudest lately. Five numbers, tracked weekly, per site: pizzas per labour hour (total pizzas produced divided by total kitchen labour hours that week — your core productivity number), dough waste percentage (from Chapter 2, the share of mixed dough that never became a sold pizza), average delivery time from order confirmation to door, cheese cost as a percentage of pizza revenue (your best single early-warning signal for both topping creep and supplier price movement), and average review score across your main platforms for the trailing 30 days.

Each number matters less in isolation than it does in the pattern it forms alongside the other four, and that's the actual skill in reading the scorecard well. A site with rising cheese cost percentage and stable pizzas-per-labour-hour is probably experiencing topping creep or a supplier price change — check the Slice Margin Matrix and the Supplier Scorecard before assuming anything else. A site with falling pizzas-per-labour-hour and stable dough waste is probably overstaffed for its current volume, or running an inefficient shift pattern, not a quality problem at all. A site with rising delivery time and falling review score together, with everything else stable, points straight at the density and radius math from Chapter 7 rather than anything happening inside the kitchen.

Illustratively, a four-site business running the scorecard might see: Site A, pizzas/labour hour 14.2 (target 13.5), dough waste

1.8%, delivery time 21 minutes, cheese cost 19.1% of pizza revenue, review score 4.6 — clearly the strongest site, worth studying for what it's doing right rather than just leaving alone. Site D, pizzas/labour hour 10.1, dough waste 4.2%, delivery time 29 minutes, cheese cost 23.4%, review score 4.1 — behind on every single metric, which tells you this isn't one isolated problem but a site that needs a genuine intervention, probably starting with a direct conversation with the manager using the weekly review format from Chapter 10, working through the Manager Mandate to check whether standards have drifted or whether the manager is under-resourced.

What to do when one site is behind matters as much as spotting it. The instinct is often to swoop in and fix everything at once, which overwhelms a struggling manager and rarely produces a lasting improvement. Better: pick the single worst-performing metric relative to your best site, fix that one thing first with a clear, time-boxed plan — say, two weeks to bring dough waste from 4.2% down toward 2.5% using the Fermentation Clock from Chapter 2 — verify the fix held for two consecutive weeks, then move to the next metric. A struggling site rarely has one root cause; it usually has three or four compounding problems, and trying to fix all of them simultaneously is how a manager ends up more overwhelmed than before you tried to help.

The scorecard also protects your best sites, which is a benefit owners underrate. Without a regular, comparable measurement, a strong site's slow decline can go unnoticed for months, because "still doing fine relative to memory of how it used to be" isn't the same question as "still doing fine relative to this week's

actual number." A site drifting from 14.2 pizzas per labour hour down to 12.8 over four months is a meaningful decline that a monthly glance at the bank balance will never surface in time to address it while it's still a small, easy fix.

Keep the scorecard visible, not private. A version of it shared with every site manager — showing all sites side by side, not just their own — tends to do more for standards than any memo ever will, because most managers running a good site take real pride in seeing it, and most managers running a behind site would rather fix it quietly before the next week's numbers post than have a private conversation about why they're behind. I've watched this transparency alone close a meaningful part of a gap between a business's best and worst site within two or three months, purely because nobody wants to be visibly last on a list their peers can see.

This week: set up the five-metric scorecard for every site you currently operate, pull the last four weeks of real data for each metric even if some of it takes digging out of your POS and delivery platform reports, and identify which single site and which single metric needs attention first.

CHAPTER 15

Financing Site Two To Five

The owners who get hurt worst in this business are rarely the ones who grow too slowly. They're the ones who grow too fast on capital that was never really theirs to risk, borrowing against a good first site's cash flow to fund a second and third opening on a timeline set by ambition rather than by what the business itself could actually sustain. A pizzeria concept that would have thrived opening its second site eighteen months later, with more cash cushion and a better-tested Manager Mandate, instead opens on a tight loan schedule, and the debt service on that loan becomes a permanent tax on every site's margin for years, turning what should have been a strong, growing business into one that's perpetually one bad month away from a genuine crisis.

The starting discipline is understanding what a new pizzeria site actually needs to prove before it's carrying its own weight, and holding every new site to the same payback standard regardless of how confident you feel about it. That's **the site investment model**: total opening investment (fit-out, equipment including your oven, initial inventory, deposits, pre-opening staff training costs) divided by the site's projected monthly free cash flow once it reaches stable operation — not its opening-month numbers, which are almost always distorted by a launch promotion and curiosity traffic that doesn't represent steady-state demand. That division gives you a payback period in months, and it's the single

number every new site proposal should be judged against before you sign a lease, not after.

Illustratively: a new site costs €165,000 to open fully fitted, including a deck oven, and is projected, based on the Pizza Site Score from Chapter 9 and comparable performance at your existing sites with a similar density and competition profile, to generate €7,200 a month in free cash flow once stable, typically reached by month four or five. That's a payback period of roughly 23 months. Compare that against a rule I give owners explicitly: **the Payback Rule of 24** — a new pizzeria site should be projected to pay back its full opening investment within 24 months of reaching stable operation, and if your honest projection runs meaningfully longer than that, the site either needs a lower opening cost, a stronger location, or simply more patience before you open it, rather than being pushed forward anyway on optimism about how quickly it'll outperform projection.

The rule exists because pizza specifically has a wider range of realistic outcomes than owners assume going in — a site can underperform its density-based projection by 20–30% in its first year simply because building neighbourhood awareness and repeat-customer habit for a pizza brand takes longer than owners expect, especially before the direct-ordering migration from Chapter 6 has had time to work. A 24-month payback target, rather than a more aggressive 12- or 15-month target some growth-hungry owners reach for, builds in enough margin to survive a slower-than-projected first year without the site's debt service becoming an emergency.

On financing sources specifically: cash flow from existing, stable sites is the cheapest and lowest-risk source of growth capital you have, and it should fund as much of each new site as your patience allows, because it carries no interest and no covenant that can get tightened if a quarter runs soft. Bank lending, when needed to bridge the rest, should be sized so that debt service on the new site's own projected cash flow — not your combined business's cash flow — comfortably covers the loan payment with margin to spare, because a loan that only works if every site performs to plan simultaneously is a loan that will eventually fail during exactly the ordinary rough quarter every business has sooner or later. Outside investors, bringing in equity partners specifically to fund expansion, are worth their own careful chapter later in this book (Chapter 19 covers the trade-offs in full), but the short version here: equity capital dilutes your ownership and often your control over exactly the local, hands-on decisions that made your first site work, and it should be a deliberate strategic choice for accelerating growth you've already proven works, not a way to paper over a business that can't yet fund its own expansion from cash flow and reasonable debt.

The growth pace rule I give owners as a final guardrail, on top of the per-site payback discipline: don't open a new site until your most recently opened site has been stable — meeting its own projected numbers on the Multi-Site Scorecard from Chapter 14 — for at least three consecutive months. This sounds conservative, and it is deliberately conservative, because the temptation to open site four while site three is still finding its feet is strong when growth momentum feels good, and it's exactly

how an owner ends up managing two unstable sites simultaneously instead of one, with half the attention available for each.

I've watched the growth pace rule get broken most often not out of greed but out of fear — a good lease becomes available in a location the owner has wanted for years, and the worry that it'll be gone if they wait three more months overrides the discipline. Sometimes that fear is justified and the site really is a rare opportunity worth an exception. More often, a genuinely good location for a pizzeria isn't a once-in-a-lifetime event — good density, reasonable competition, and adequate utility capacity exist in more places than owners assume once they've built the Pizza Site Score habit from Chapter 9, and the next good lease is usually less rare than it feels in the moment. Treat an exception to the pace rule as a real decision requiring real justification, not a default response to a nice-looking listing.

This week: calculate the real payback period for your most recent site opening using actual numbers rather than the original projection, and compare it honestly against the Payback Rule of 24 — if it's running long, use that as real information about what your next site's opening cost or location needs to look like, not as a reason to just hope the next one performs better.

CHAPTER 16

The Slice Window, The Kiosk And Other Formats

By the time an owner has three or four full-service or delivery-and-pickup sites running well, a different kind of opportunity usually starts presenting itself: a small slice-counter unit in a food hall, a market stall opportunity at a weekend event, a dark-kitchen space with cheap rent and no dine-in requirement at all. These smaller formats are genuinely tempting, because they promise brand extension at a fraction of a full site's opening cost — and they can be excellent additions to a multi-site pizza business, or they can quietly dilute the brand and drain management attention from the sites actually carrying your revenue, depending entirely on whether you evaluate them with the same rigour you'd apply to a full site.

They also tend to arrive at exactly the moment an owner is feeling most confident, which is worth noticing in itself. Three or four sites in, the Multi-Site Scorecard is looking healthy, the Manager Mandate is working, and a small-format opportunity shows up dressed as a low-risk way to keep the momentum going without the full weight of another complete site opening. That framing — low capital, therefore low risk — is exactly the trap this chapter exists to name, because capital risk and attention risk are not the same risk, and a format that's cheap to open can still be expensive in the currency that actually matters

most at this stage of your growth: your own and your best managers' limited hours.

The mistake I see most is treating a small format as a low-stakes decision simply because the capital outlay is low. A slice window costing €18,000 to fit out feels like a rounding error next to a €165,000 full site, so owners skip the Pizza Site Score, skip a real payback calculation, and greenlight it on enthusiasm — and then discover that a small format still consumes real management attention, real staff scheduling complexity, and real brand-consistency risk, none of which scales down proportionally just because the rent did.

The format comparison matrix forces the same discipline onto every format under consideration, scored on four dimensions that matter differently for a kiosk than for a full site. Brand fit: does this format let you deliver your actual product — your real dough, baked properly — or does the format's constraints (a slice window reheating pre-baked slices rather than baking fresh, for instance) mean customers are experiencing a lesser version of your brand that could damage trust in the full version. Margin per labour hour: small formats often look attractive on rent but can be brutal on labour efficiency if a single staff member is required regardless of how thin the actual order volume runs, meaning your labour cost as a percentage of revenue can be structurally worse than a full site even with lower absolute costs. Brand visibility value: what does this format do for awareness of your full sites nearby, separate from its own direct profitability — a market stall near a planned fourth full-site location can be a legitimate customer-acquisition investment even at thin or

negative direct margin, as long as you're honest that's what it is rather than pretending it as a standalone profit centre. And operational complexity added: does this format require a genuinely different skill set, different equipment, or different supply chain that adds real complexity to your Chapter 12 purchasing consolidation and Chapter 11 training ladder, or does it plug cleanly into systems you already run.

Illustratively, run three format candidates through the matrix. A slice window in a food hall, reheating pre-portioned slices from your commissary: strong on margin per labour hour (one person, high throughput, low skill requirement), weak on brand fit (a reheated slice is a genuinely different product experience than a fresh bake, and customers who try only this format may never discover what your full sites actually offer), moderate on visibility, low on added complexity since it draws on commissary supply you already run. A dark kitchen, full menu, delivery-only, no customer-facing space at all: strong on margin per labour hour if delivery density supports it (Chapter 7's math applies in full here), zero brand-visibility value since there's no physical presence for anyone to discover, low on brand-fit risk since it's your actual full product, moderate complexity since it still needs a trained pizzaiolo team. A weekend market stall using a portable oven: weak on margin per labour hour (labour-intensive relative to thin, weather-dependent volume), strong on brand visibility especially in a target neighbourhood for a future full site, genuinely good brand fit if you commit to baking fresh on-site rather than reheating, and meaningfully higher operational

complexity than either of the other two because of the portable equipment and event logistics involved.

None of these three is universally right or wrong — the framework's value is forcing you to be honest about which dimension a given format is actually winning on, so you're not accidentally justifying a weak-margin format with an unexamined assumption about brand visibility it isn't actually delivering, or accepting real brand-fit risk from a reheated-slice concept because the headline margin number looked good in isolation.

The financial test I'd add on top of the matrix: never let a small format's opening cost trick you into skipping the Payback Rule of 24 from Chapter 15. A €18,000 slice window still needs to pay itself back within a reasonable window against its own realistic free cash flow, and "it's cheap, so who cares if it's slow to pay back" is exactly the reasoning that leads to three small formats quietly draining management time for years while never becoming genuinely profitable additions to the business.

One last consideration before you commit to any small format: think about the exit as carefully as the entry. A full site's lease and equipment investment forces a certain discipline around closing it down if it isn't working, precisely because the stakes are high enough that the decision gets made deliberately. A small format's low stakes cut both ways — easy to open, but also easy to let limp along indefinitely at break-even or worse, because no single week's loss is ever quite bad enough to trigger the hard conversation about shutting it down. Set a review date and a

clear numeric threshold when you open it, not just when you're deciding whether to open it.

This week: if you're considering, or already running, a smaller-format extension, score it honestly on the four-dimension format comparison matrix, and calculate its real payback period the same way you would a full site — a small commitment still deserves a real analysis, not an exemption from one.

CHAPTER 17

Reputation Across Sites

Here is a fact about how customers actually search for pizza that changes the reputation calculus completely once you have more than one site: they are not searching for "your brand name." They're searching for "pizza near me," and the search results and map listings they see are frequently aggregated or influenced by your brand's overall rating and review signal across every location, not just the one nearest them. This means a badly managed site three's one-star reviews don't just hurt site three — depending on the platform and how your listings are structured, they can quietly suppress how prominently your genuinely excellent site one shows up to a customer who has never set foot in site three and never will.

I've had owners push back on this the first time they hear it, because it feels counterintuitive that a stranger's bad night at a location across town could affect their own excellent local site's visibility. The mechanism varies by platform and changes over time as ranking algorithms evolve, so I won't pretend to know the exact current weighting any given platform applies — but the directional pattern has held consistently across the businesses I've watched for years: brand-level aggregate signal matters more than any single owner expects it to, and it's a risk that simply didn't exist when there was only one site to manage.

I call this **the Weakest Site Rule**: in a reputation sense, your multi-site business is only as strong as its worst-performing

location, because that's the location dragging down your aggregate signal, generating the review pattern that shows up when a prospective customer scans your brand's overall profile before deciding whether to try any of your locations at all. A single-site owner never faces this dynamic — their reputation is simply their reputation. A multi-site owner is, whether they've thought about it this way or not, managing a shared reputation asset that every site both draws from and contributes to.

This makes site-level review management a genuinely different job than it was at one location, and it needs its own routine rather than being left to whichever manager happens to think of it. **The multi-site reputation routine** runs on three cadences. Daily: every new review across every platform, at every site, gets read and responded to within 24 hours by the site manager, using a response that actually addresses the specific complaint or compliment rather than a generic copy-pasted reply — customers reading your review section can tell the difference immediately, and a pattern of generic replies reads as not caring even when individual complaints are being handled fine operationally. Weekly: review scores and review volume per site get added as a line to the Multi-Site Scorecard from Chapter 14, so a declining trend at any site surfaces in the same regular rhythm as your other operational numbers rather than being noticed only when it's already become a crisis. Monthly: someone — the owner, or a designated person if you're past the size where the owner can do this alone — audits every site's platform listings directly: correct hours, correct address and phone number, correct menu with current prices, correct photos

that actually represent your current product rather than photos from an opening three years ago that no longer match what leaves the kitchen.

That listing-hygiene piece sounds trivial and isn't. I've seen sites lose meaningful delivery-platform order volume purely because a listed closing time was wrong after a schedule change, silently turning away customers who tried to order during hours the platform incorrectly showed as closed — a completely invisible loss, because a customer who sees "closed" simply orders from a competitor instead and never tells you why. Wrong photos cause a subtler version of the same problem: a customer who orders based on an appealing photo and receives a pizza that looks meaningfully different is primed to feel disappointed before they've even opened the box, independent of whether the actual pizza is good.

Handling a genuinely bad review — the kind that's specific, credible, and damaging, not a clearly bad-faith one-star with no detail — deserves its own approach, because how you respond publicly matters as much to future customers reading it as it does to the original complainant. Acknowledge the specific issue without being defensive, state plainly and specifically what you're doing about it (not a vague "we take this seriously," but the actual fix — "we've adjusted our packaging as of this month specifically to address the issue you described"), and take the resolution offer offline via a phone number or email rather than negotiating a refund in the public comment thread. A well-handled bad review, visible to every future customer who reads it, can do more for your credibility than an unbroken string of five-star

reviews with no evidence you handle problems well when they occur — because every customer knows problems occasionally happen, and what they're actually evaluating is how you respond when they do.

Illustratively, a four-site business tracking review score as part of its scorecard might notice Site C's average dip from 4.5 to 4.1 over six weeks, with the dip concentrated in complaints mentioning delivery time specifically. Cross-referencing against the same period's delivery-time metric on the scorecard confirms it: Site C's average delivery time had crept from 24 to 33 minutes over the same window, likely tied to a new driver still learning the route density from Chapter 7. The reputation dip wasn't a standalone problem needing a reputation-specific fix — it was a downstream symptom of an operational issue the scorecard would have caught anyway, but the review pattern caught it faster and gave a clearer, more specific signal of exactly what customers were experiencing.

Assign the daily response duty to the site manager, not to you centrally, even though a template or two of guidance from you is worth providing. A response written by the manager who actually runs that kitchen carries a specificity and authenticity a centrally written reply never will, and it keeps the manager directly connected to what customers are actually experiencing at their own site — which is, not incidentally, one more channel through which a struggling site's problems surface early, alongside the Multi-Site Scorecard from Chapter 14.

This week: set up the daily-weekly-monthly reputation routine formally if you haven't already, starting with a monthly listing-hygiene audit across every platform and every site today — it's the fastest of the three to do once and the easiest to have quietly gone stale without anyone noticing.

CHAPTER 18

Central Systems Without A Head Office

There's a specific fear I hear from owners approaching their fourth or fifth site, and it's worth naming directly because it stops good decisions from getting made: the fear of turning into a head-office business. Owners who got into this to make good pizza and run a good local business watch chains build layers of regional managers, corporate systems, and administrative overhead, and they recoil from it — reasonably, because that overhead is exactly what makes a chain feel like a chain rather than a collection of businesses a neighbourhood still considers theirs. The mistake is concluding from that reasonable fear that no central systems at all is the right answer, when the actual right answer is a small number of genuinely load-bearing shared systems and nothing more.

The irony worth sitting with is that avoiding all central systems out of fear of becoming a chain often produces the exact outcome the fear was trying to prevent, just by a different route. Without a shared recipe system, sites drift apart on dough and spec until the brand genuinely does fragment into loosely related local businesses — the local-irrelevance trap from Chapter 13, arrived at not through deliberate local adaptation but through simple neglect. A handful of the right shared systems is what actually protects local character, by handling the boring consistency work centrally so each site manager's own time and judgment can go toward the things that genuinely benefit from

being local: how they talk to their own regulars, how they run their own kitchen's daily rhythm, how they represent the brand to their own neighbourhood.

The test for what belongs centrally isn't "would this be more efficient centralised" — almost everything passes that test in theory. The test is: does this system fail specifically, expensively, and repeatedly at every site if it isn't shared. That's a much narrower set, and **the Five-System Rule** names it explicitly, because five is roughly how many genuinely load-bearing shared systems a five-site independent pizza business actually needs, not the dozen-plus a chain's head office typically runs.

System one: a shared point-of-sale platform across every site, because it's what makes the Multi-Site Scorecard from Chapter 14 possible at all — comparing sites on pizzas-per-labour-hour or cheese cost percentage requires the underlying sales and labour data to be structured identically, which a site running its own independently chosen POS simply can't provide. System two: a shared recipe and specification system — your Dough DNA Sheet, your Slice Margin Matrix specs, your Manager Mandate — held in one place every site references, rather than living as a folder of documents that quietly diverge in local copies over time as each site's manager makes their own small edits. System three: a shared roster and scheduling tool, because labour is your second-largest cost after food, and a shared system lets you see labour efficiency patterns across sites (feeding directly into the scorecard) rather than trusting each site's own spreadsheet or paper rota. System four: one bookkeeping process — not necessarily one bookkeeper, but

one consistent chart of accounts and one consistent monthly close process across every site, because comparing site profitability meaningfully requires the underlying numbers to be built the same way everywhere, and I've watched owners discover, only after several years, that two of their sites had been categorising costs differently enough that a site they believed was their strongest performer actually wasn't once the numbers were normalised. System five: a single supplier and purchasing record, tying back to Chapter 12's consolidation plan and Single-Source Risk Score, so purchasing decisions and supplier performance history live in one place rather than in each site manager's personal relationship with their local rep.

Everything else — and this is the part that protects you from becoming a head-office business by accident — stays local, deliberately. Hiring decisions for line staff, day-to-day supplier ordering within agreed par levels, local marketing execution, the specific rhythm of a site's daily prep schedule: these belong to the site manager, inside the Manager Mandate's delegated list from Chapter 10, and centralising them wouldn't make the business better, it would just make it slower and less responsive to the actual conditions each site manager can see and you, working across five locations, structurally cannot.

Build **the systems map and cost budget** by listing your five systems, what each currently costs monthly (software subscriptions, the portion of anyone's role dedicated to running it), and what specifically breaks at which site count if that system isn't shared — this is the discipline that keeps the list at five rather than letting it creep toward twelve, because every

proposed addition has to clear the same bar: what specifically, expensively, and repeatedly fails without it. Illustratively, a five-site business's systems map might show: shared POS at €340/month total, recipe and spec system essentially free (a well-organised shared document library plus discipline), roster tool at €95/month, bookkeeping process costing roughly six hours a month of a part-time bookkeeper's time per site, and purchasing records folded into the same document system as recipes. Total genuinely central-system cost: often surprisingly low, well under what a single junior head-office hire would cost, which is precisely the point — the systems, not the staff, do most of the coordinating work a chain would otherwise pay a layer of management to do.

Revisit the list annually rather than treating it as fixed once decided, because the right answer to "what needs to be shared" genuinely shifts as your site count grows. A system that felt unnecessary at three sites — a formal purchasing record, say, when personal relationships with a handful of suppliers were enough to manage informally — can become genuinely load-bearing at seven or eight, once the number of relationships and the volume of orders exceeds what any one person can track from memory. The Five-System Rule isn't a permanent ceiling; it's a discipline for asking the right question at whatever size you currently are, and the honest answer to that question changes as you grow, even if the number of systems it points to stays close to five for longer than owners expect.

This week: list your own five candidate central systems against the "what specifically breaks without this, at every site,

repeatedly" test, and be honest about anything currently centralised that doesn't actually clear that bar — it's probably slowing your managers down for no real benefit.

CHAPTER 19

Franchise, Partner Or Own

Past five sites, growth stops being a question you can answer with better systems alone, and becomes a question about capital structure and control that most operators have never had to think through before, because nothing about running one excellent pizzeria prepares you for deciding how to fund and govern site six through twenty.

This is also the point where well-meaning advisors — an accountant, a lawyer, a friend who scaled a different kind of business entirely — start offering opinions with real confidence and very little grounding in what makes pizza specifically different from the retail or services businesses their experience is drawn from. Franchising advice built around a business where execution quality barely varies between operators does not transfer cleanly to a business where the gap between a competent and an excellent pizzaiolo is the difference between a five-star and a two-star customer experience. Take outside advice on capital structure and legal mechanics gladly; be far more skeptical of outside advice on how much quality-control risk your specific product can tolerate under someone else's ownership.

Three broad models exist, and each trades control for capital and pace in a different way. Company-owned growth — every new site funded and operated exactly as sites one through five were, under the Payback Rule of 24 and Growth Pace Rule from

Chapter 15 — keeps full control and full economics in your hands, but is capital-constrained by definition: you can only grow as fast as your own cash flow, and reasonable debt against it, allows. Operating partners with equity — bringing in a trusted manager, often someone who's proven themselves running one of your existing sites well, as a genuine minority equity owner of a new site rather than just a salaried manager — solves your capital constraint partially (they typically invest some capital alongside you) and solves a deeper problem company-owned growth doesn't: an equity-owning operating partner has a fundamentally different, more durable motivation to protect quality and margin at their specific site than an employed manager does, because it's now genuinely their business too, not just their job. Franchising — licensing your brand, systems, and training to an independent operator who funds and legally owns their own site under your brand standards — solves the capital constraint completely, since franchisees fund their own openings, but costs you the most control of the three models, and control is precisely the thing this entire book has been about protecting.

The Control-Capital Grid maps these three models against the two things that actually matter for a pizza-specific decision: how much capital each model requires from you personally, and how much direct control you retain over the exact quality standards — the Dough DNA Sheet, the Manager Mandate, the Slice Margin Matrix — that this whole book argues are the actual product. Company-owned sits at maximum control, maximum capital requirement. Operating-partner equity sits at high control (you're

still setting and enforcing every core standard, just sharing upside with someone who has real skin in the game) and moderate capital requirement. Franchising sits at the lowest capital requirement and, for pizza specifically, the lowest control — and this is the trade-off that deserves the most scrutiny in this chapter, because pizza is unusually control-sensitive compared to many franchised concepts.

Here's why pizza specifically punishes weak control harder than many other franchisable concepts do. A franchised sandwich or coffee concept can survive meaningful variation in execution skill across franchisees, because the product's quality ceiling and floor are both relatively close together — a mediocre execution and an excellent one aren't wildly different experiences. Pizza's quality range between a franchisee who nails the Dough DNA Sheet and one who treats it as a loose guideline is enormous, because fermentation, oven skill, and peel technique are genuinely difficult crafts with a wide gap between competent and excellent — and a customer who has a bad experience at any franchised location under your brand name attributes that experience to your brand as a whole, exactly as described in Chapter 17's Weakest Site Rule, except now the weak site isn't even one you directly control.

This is why the operators I've watched succeed longest at growing past five sites while staying genuinely independent in spirit tend to lean toward operating partners with equity over pure franchising, specifically because of pizza's unusual sensitivity to execution quality. An equity partner is still trained through your Pizzaiolo Ladder, still held to your Manager Mandate, still

measured on your Multi-Site Scorecard — the control mechanisms from earlier chapters remain fully in force — while gaining a capital contribution and a motivation structure closer to an owner's than an employee's. Franchising can still work for pizza, and does for plenty of large, successful brands, but it typically requires a level of training infrastructure, quality-audit staffing, and enforcement willingness — including the willingness to actually terminate a franchisee's agreement over a real quality failure, not just issue warnings — that a genuinely independent operator moving past five sites usually isn't ready to build yet, and often shouldn't rush to build before they've proven the equity-partner model first at sites six through eight or nine.

None of the three models is permanent, either, and the sequencing matters more than the final choice. Most owners I've watched navigate this well started with company-owned sites through five or six locations to prove the systems in this book actually hold at that scale, moved to operating-partner equity for the next several as capital and management-attention constraints started to bind, and only considered franchising, if at all, once the training and quality-audit infrastructure genuinely existed rather than being aspirational. Skipping straight to franchising because it's the fastest-capital option, before proving the model holds under partners who are still directly accountable to your own standards, is the sequencing mistake that costs the most in pizza specifically, precisely because of how much execution quality varies and how visibly customers notice.

This week: if you're approaching or past five sites, map your own growth ambition against the Control-Capital Grid honestly, and if

franchising is on your mind, ask yourself specifically whether you currently have the training and quality-audit infrastructure this chapter describes — not whether the concept could theoretically franchise, but whether you personally could enforce your own standards on someone who owns their site independently of you.

CHAPTER 20

Staying Independent At Scale

I want to end where every chapter in this book has quietly been pointing: the goal was never simply "more sites." More sites is easy to want and easy to get wrong in exactly the ways this book has spent nineteen chapters describing. The actual goal, the one worth the years of work multi-site growth demands, is more specific and harder to hold onto: multiple sites that still taste like the first one, an owner who is not personally baking every night just to keep quality from slipping, and a brand that the people in each neighbourhood still consider genuinely theirs rather than a chain that happened to open nearby.

If there's one idea I'd want you to carry out of this book above every named framework in it, it's **the First-Site Test**: at any point in your growth, at any site count, ask honestly whether a regular customer from your very first location, transplanted to your newest one, would recognise what they're eating and how they're being treated. Not identical in every superficial way — Chapter 13 already made the case for genuine local flex — but recognisable, unmistakably, as the same business they fell in love with originally. Every tool in this book exists to help that test keep passing as the distance between you and any given site grows. The day it stops passing quietly, without anyone deciding it should, is the day a business that was genuinely trying to stay independent at scale has, without quite meaning to, become something else.

Every framework in this book serves that one goal from a different angle, and it's worth seeing them together as a single system rather than twenty separate tools. The Copyability Scorecard and Dough DNA Sheet make your first site's excellence transferable in the first place. The Commissary Break-Even and Peel-to-Plate Rate make sure your production model can actually support the volume you're asking it to support. The Slice Margin Matrix and Channel P&L Per Order protect the margin that funds everything else. The Pizza Site Score and Density-Radius Model make sure you're building on genuinely sound ground rather than an exciting-looking mistake. The Manager Mandate, the Pizzaiolo Ladder, and the Multi-Site Scorecard let you trust people who are not you, at a standard you can actually verify rather than just hope for. The Core-and-Flex Rule and the Weakest Site Rule protect what the brand means across every neighbourhood it touches. The Payback Rule of 24, the Growth Pace Rule, and the Control-Capital Grid make sure you're growing at a pace and through a structure your actual capital and actual control tolerance can sustain. None of these tools works in isolation nearly as well as it works alongside the others — a strong Dough DNA Sheet doesn't protect you from opening the wrong site, and a great site selection doesn't protect you from a manager operating with no defined mandate.

I want to close with a five-year planning exercise, because a book full of weekly and monthly tools should also help you see further out than next week. Build **a five-year plan** across three tracks, reviewed and adjusted annually rather than treated as fixed the day you write it. The quality track: your target

Copyability Scorecard and Pizzaiolo Ladder certification levels across every site, checked against your actual measured numbers each year rather than assumed. The profit track: your target site count, your target payback period per new site staying disciplined to the Payback Rule of 24, and your target Multi-Site Scorecard performance band — not every site needs to hit your best site's numbers, but every site needs to sit within a defined, acceptable range of them. The owner-time track, and I'd argue this is the one most books skip entirely and the one that actually determines whether you'll still love this business in five years: how many hours a week you're spending on which of the three roles a multi-site owner actually plays — hands-on production and training, management and coaching of your site managers, and strategic decisions about sites, systems, and financing — and how you want that split to shift as you add sites, so growth is something you're deliberately shaping rather than something that's quietly reshaping you.

Illustratively, a five-year plan for an owner currently at three sites might target: eight sites by year five, average Copyability Scorecard of 20+ across all locations (up from a current honest average of 15), Pizzaiolo Ladder coverage of at least two stage-four-certified pizzaiolos per site, payback periods held under 22 months on every new opening, and — the owner-time track — a deliberate shift from roughly 60% hands-on/25% management/15% strategic time today toward 15% hands-on/40% management/45% strategic by year five, with hands-on time preserved specifically for training new pizzaiolos through the ladder rather than covering shifts, because that's the hands-on

work that actually compounds rather than the work that just fills a gap for one night.

The owners I've watched build something genuinely worth having at five, six, seven sites share one trait more than any other: they never stopped treating their first site's dough, their first site's standards, as the actual definition of the business, no matter how many locations, managers, or systems got added on top of it. Growth, done well, is the discipline of protecting that definition at increasing distance from yourself — not abandoning it for something bigger, and not staying so small that the thing you built never gets to reach the people in the next neighbourhood who would genuinely love it if they ever got the chance to taste it.

This week, and really for the coming year: write your own five-year plan across the quality, profit, and owner-time tracks, with real numbers on each, and put a date on your calendar six months from now to check your actual figures against it — a plan nobody revisits is just a nice document, and the businesses in this book that actually got where they wanted went were run by owners who treated their plan as a working document, not a keepsake.

Glossary

Copyability Scorecard	An honest rating of how much of your first site's success actually lives in written systems versus in your own daily presence.
Dough DNA Sheet	The written specification for your dough: hydration, fermentation windows and cold-proof logistics, kept identical across every site.
Commissary Break-Even	The site-count calculation showing when central dough production becomes cheaper than making dough fresh at each location.
Peel-to-Plate Rate	Pizzas-per-hour through the oven, the hard capacity limit of a site regardless of how many tables or drivers it has.
Slice Margin Matrix	A per-pizza margin table used to keep the menu short and protect profit as a chain grows across sites.
Channel P&L Per Order	The true profit of one order broken out by fulfilment channel (platform delivery, own delivery, pickup), after commission and packaging.

Density-Radius Model	A method for setting delivery radius by order density per postcode rather than by raw distance from the site.
Fifteen-Minute Box Test	A packaging test that checks whether a pizza still arrives correctly after the typical delivery travel time.
Pizza Site Score	A scoring model for a candidate location covering residential density, footfall, competition and utility supply.
Manager Mandate	The written definition of what a site manager may decide alone, and what still needs the owner, at each stage of growth.
Pizzaiolo Ladder	A skills path from stretch to bake to peel, with time and quality standards at each rung, used to train without depending on one star cook.
Single-Source Risk Score	A rating of how exposed the business is when one ingredient (commonly the mozzarella) has only one supplier.
Core-and-Flex Rule	The principle that a brand keeps a small, fixed core across every site while allowing a defined, limited local flex.

Multi-Site Scorecard	The weekly comparison sheet (pizzas per labour hour, dough waste, delivery time, cheese cost, review score) used to compare sites fairly.
Payback Rule of 24	A rule of thumb for how many months a new site's investment should take to pay back before the next site is opened.
Weakest Site Rule	The discipline of fixing the worst-performing site before opening a new one, since one bad site drags the whole brand's ratings.
Five-System Rule	The minimum shared systems (POS, recipes, roster, purchasing, bookkeeping) a five-site independent needs without building a head office.
Control-Capital Grid	A framework for weighing growth models (company-owned, partner-operator, franchise) against how much control and capital each keeps.
First-Site Test	The question asked before any expansion: would this new site still taste and feel like the first one a regular already loves?

About the Author

Thibault Van de Sompele is the founder of HappyChef, a reservation and operations platform used by independent restaurants and hotels across Europe. Before building software, he spent years working alongside hospitality operators directly, and it was that closeness — sitting with owners as they tried to make sense of their own booking patterns, labour costs, and kitchen bottlenecks — that convinced him the real gap in the market wasn't another generic booking tool but a way of seeing a business's own numbers clearly enough to act on them.

HappyChef now sits inside the daily operations of hundreds of independent hospitality businesses, from single-site cafés to multi-location groups managing several kitchens at once. That vantage point, watching the same patterns of success and struggle recur across very different concepts, cuisines, and countries, is what shaped this book. Pizza businesses in particular kept surfacing a specific, recognisable pattern: brilliant single sites, and a second location that struggled in entirely predictable ways nobody had written down clearly before.

Thibault lives and works in Essen, Belgium, where HappyChef is based, and remains closely involved with the operators who use the platform day to day. He writes for owners who are building something they intend to still be proud of at five locations, not just at one — businesses that grow without losing the thing that made them worth growing in the first place.