

THE SHIFT UP: OPEN

The Complete Launch Guide
for Restaurants, Bars,
and Hospitality Venues



THE COMPLETE PLAYBOOK TO PLAN,
LAUNCH & PROFIT FROM DAY ONE

MARCUS TREAMER

Foreword by Steve Cromie

The Shift Up: Open

The Complete Launch Guide for Restaurants, Bars,
and Hospitality Venues

Marcus Treamer

www.maketheshiftup.com

Copyright © 2026 Marcus Treamer. All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without prior written permission, except for brief quotations used in reviews or critical commentary.

This book is a practical guide for hospitality operators. It is not legal, financial, tax, or regulatory advice. Always consult qualified local professionals before making binding decisions.

Contents

Foreword	1
Introduction	5
Part One: The Concept	9
Chapter 1: Finding Your Concept	11
Chapter 2: Know Your Market	23
Chapter 3: Building Your Brand	35
Chapter 4: The Concept Document	45
Chapter 5: Concept Validation	55
Part Two: The Plan	65
Chapter 6: The Business Plan	67
Chapter 7: Capital Requirements and Runway	77
Chapter 8: Raising Investment	85
Chapter 9: Legal Foundations	95
Chapter 10: The Project Plan	105
Part Three: The Venue	113
Chapter 11: Finding Your Location	115
Chapter 12: Design, Layout, and Build-Out	123
Chapter 13: Licensing, Permits, and Compliance	133
Chapter 14: Building Your Team	141
Chapter 15: Menu Design, Beverage Programme, and Suppliers	149
Part Four: The Opening	157
Chapter 16: Pre-Opening Marketing	159
Chapter 17: Your Digital Presence	167
Chapter 18: Agencies, Influencers, and Partnerships	175
Chapter 19: The Soft Launch	183
Chapter 20: Opening Night and the First 90 Days	191
Conclusion	203
Resource Directory	206
Glossary	207
Worksheet Index	209
Recommended Reading	210
About the Author	211

Foreword

By Steve Cromie

I have been in this industry for over thirty years. I started pulling beers and spinning records on Australia's east coast when I was barely old enough to drink what I was pouring. I have worked my way from bartender to COO across four continents, opened venues from concrete slabs, turned around operations that were weeks from closing, and built the world's highest bar on the 122nd floor of the Burj Khalifa. Along the way, I have made every mistake this book warns you about – and a few it was too polite to mention.

I tell you this not to impress you. I tell you because when I read what Marcus has written, I kept thinking the same thing: *where was this book when I needed it?*

The industry I entered at eighteen did not have guides like this. You learned by doing. You learned by getting it wrong. You learned because the head chef threw a pan at you – not because someone sat you down and explained how cost of goods sold works or why your concept document matters more than your mood board. The education was brutal, expensive, and entirely dependent on whether the people above you were generous enough to share what they knew. Many were not. They hoarded knowledge because they thought it protected their position. It did not. It just meant the next generation had to re-learn the

same lessons from scratch, at full price.

Marcus is not that kind of operator. He shares everything. Aggressively, methodically, and with the kind of obsessive detail that makes you wonder whether he has ever slept a full night in his life. I have known him for years, worked alongside him on projects that would test anyone's patience, and watched him do the one thing that separates the operators who last from the ones who flame out: he builds systems. Not because systems are glamorous – they are the opposite of glamorous – but because systems are what keep a venue alive at eleven o'clock on a Saturday night when everything that can go wrong is going wrong simultaneously and the only thing standing between you and disaster is the process you put in place three months ago.

This book is a system. It is the system for opening a hospitality venue – not the romanticized version where everything falls into place because the food is good and the dream is pure, but the real version where the contractor sweeps concrete into your drain, the alcohol licence arrives three months late, and your best hire quits during the soft launch because they got a better offer from the place down the road. Marcus has mapped every one of those scenarios, not from theory but from having lived them – in Hong Kong, Shanghai, Dubai, Beijing, Koh Samui, and everywhere in between.

I contributed a few stories to this book. The Burj Khalifa, where a world-renowned designer handed me a bar with no equipment inside it and told me not to ruin his work. The Armani Hotel, where the cigar lounge was measured around Giorgio Armani's frame and none of the actual guests could fit the furniture. The beach club in Koh Samui, where the architect gave me one small fridge for a two-thousand-person venue. The opening night in Patong, where a bag of concrete in a drain turned the bathrooms into a flood zone during our first event.

Every one of those stories has the same lesson: the gap between design and operations is where venues fail. Beautiful spaces that do not work when the room is full. Brilliant concepts that were never tested against the reality of a packed Friday night. This book exists in that gap. It does not tell you how to dream up a restaurant. It tells you how to build one that survives.

My father had a philosophy I have carried through my entire career: if you make a mistake once, learn from it. If you make it twice, you should have learned the first time. If you make it three times – well, his language was more colourful than I will use here, but the point was clear. This book is the first-time lesson. Every mistake documented in these pages is one you do not have to make yourself, because someone who already made it has written down what happened and what they would do differently.

Read it. Use it. And when things get hard – and they will get hard, because this industry does not do easy – remember Marcus's closing words: stay the course.

It is the best advice in this book. It might be the best advice in the industry.

Steve Cromie has over thirty years' experience in hospitality across four continents. His career includes launching the world's highest bar and restaurant at the Burj Khalifa, Dubai, and opening venues across Southeast Asia from construction through to operation. He is currently based in Southeast Asia.

Introduction

This book exists because I got tired of watching talented people lose their money, their confidence, and occasionally their marriages opening restaurants that should have worked.

Not bad people. Not stupid people. People with genuine skill, real passion, and enough capital to build something meaningful – who failed because nobody told them what they were actually getting into. Not the romantic version. Not the 'follow your passion' version. The version where the grease trap backs up during a health inspection, the investor calls on a Tuesday morning asking why revenue is 30% below projection, and the head chef you spent four months recruiting hands in their notice because someone offered them £3,000 more per year.

The hospitality industry does not lack enthusiasm. It lacks a manual.

I have spent more than twenty-five years in this industry across Hong Kong, Shanghai, Dubai, Beijing, Koh Samui, France, Spain, and Holland – opening venues, consulting on launches, and watching from close range as concepts succeeded or collapsed. The ones that collapsed rarely failed for the reasons their founders assumed. The food was usually fine. The location was usually defensible. The passion was never in question. They failed because the concept was not defined with enough precision to survive contact with the market. They failed because the capital ran out two months before the venue reached stability. They failed because the business plan projected 90% occupancy from month

THE SHIFT UP: OPEN

one and the real number was 40%. They failed because the legal structure was assembled on a handshake, the licensing timeline was underestimated by six months, or the design was built for photographs rather than service.

Every one of those failures was preventable. Not with more money. Not with more luck. With more process.

This is not a book about how to cook, how to mix cocktails, or how to design a beautiful dining room. There are excellent books on all of those subjects, and this is not one of them. This is a book about the commercial, operational, and strategic infrastructure that must be in place before the first plate leaves the kitchen – and the discipline required to maintain it during the eighteen to twenty-four months it takes for a new venue to find its stride.

The book is organized into four parts that mirror the lifecycle of opening a venue.

Part I: The Concept (Chapters 1–5) defines what you are building. You will develop a concept, research its market, build its brand, document the vision, and validate it against real evidence before a penny is committed. At the end of Part I, you will hold five documents – a Concept Canvas, a Market Research Template, a Brand Identity Brief, a One-Page Concept Document, and a Go/No-Go Decision Matrix – that together constitute a launch dossier no competitor on the shelf can match.

Part II: The Plan (Chapters 6–10) turns the concept into a business. You will write a business plan, calculate your capital requirements, raise investment on terms that protect your equity, establish the legal foundations, and build a project timeline with milestones and accountability. At the end of Part II, the concept has a commercial engine.

Part III: The Venue (Chapters 11–15) builds the physical operation. Location, design, licensing, team, and menu – the decisions that determine what your customers will actually experience. Every chapter in this section applies a single test to every decision: does this work on a Friday night at 10 PM, when the room is full and the kitchen is at capacity?

Part IV: The Opening (Chapters 16–20) brings customers through the door and keeps them coming back. Pre-opening marketing, digital presence, agency and partnership strategy, the soft launch, and the first ninety days of trading – the period that determines whether the venue reaches stability or enters a spiral.

Every chapter contains the same structural elements, because consistency in a reference book is not a stylistic choice – it is a courtesy to the reader who needs to find something at midnight when the problem is live.

Each chapter opens with a specific scenario – usually a real venue, in a real market, facing a real consequence. Each contains a minimum of five sourced data points, three to five named case studies, and at least one practical worksheet or framework that the reader can complete immediately. Each includes a Debrief in three parts (The Error, The Fallout, The Fix) drawn from my own experience or from Steve Cromie, a hospitality consultant whose career spans the Burj Khalifa, the Armani Hotel Dubai, and beach clubs across Southeast Asia. Steve's stories appear in Chapters 12, 14, 18, and 19. His fingerprints are on more of the book than the attribution suggests.

The footnotes are mine. They are the tangents, the asides, and the observations that do not fit the main argument but refused to stay in the margins. If you find them useful, that is their purpose. If you find them entertaining, that is my personality. If you find them irritating, I understand, and I wrote them anyway.

A note on geography. This book draws on markets I have worked in directly: the UK, Hong Kong, mainland China, the UAE, Thailand, France, Spain, and Holland. It also draws on research from the United States, Australia, and the broader European market. Where regulations, costs, or cultural expectations are jurisdiction-specific, the text says so. Where a principle is universal – and most of the principles in this book are universal, because the physics of running a restaurant do not change dramatically between London and Bangkok – the text presents it as such. If you are opening a venue in a market not explicitly covered, the frameworks still apply. The numbers will be different. The process will not.

A note on who this book is for. It is for the first-time operator who has the capital, the ambition, and the concept but not the roadmap. It is for

THE SHIFT UP: OPEN

the experienced hospitality professional opening their own place for the first time after years of working in someone else's. It is for the career-changer who has spent twenty years in finance or law or technology and now wants to open the wine bar they have been thinking about since their second child was born. And it is for the investor who wants to understand what a properly prepared hospitality launch looks like, so they can distinguish the operators who have done the work from the ones who have done the dreaming.

It is not for the operator who believes passion is sufficient. Passion is necessary. It is not sufficient. This book is the bridge between the two.

One final thing. Optimism is the most expensive emotion in hospitality. You will encounter that phrase more than once in the pages that follow. It is not cynicism. It is the hard-won observation that the gap between what an operator hopes will happen and what actually happens is measured in cash – and cash, unlike optimism, is finite. Every framework, every worksheet, and every data point in this book exists to close that gap. Not to eliminate optimism. To discipline it.

Let's begin.

Part One: The Concept

CHAPTER 1

Finding Your Concept – The Difference Between an Idea and a Venue

THE SHIFT UP: OPEN

In Hong Kong, on the twenty-third floor of a building overlooking Victoria Harbour, a multi-concept operator opened his fourth venue in three years. It was, by any reasonable measure, a triumph of design ambition. The interior had been featured in *Wallpaper* and *DesignBoom*. The lighting rig alone cost more than most independent restaurants spend on their entire fit-out. The cocktail bar was a curved monolith of Italian marble that photographed so well it practically captioned itself. Within six weeks, every lifestyle influencer in the city had posted from the terrace. Within eight months, the operator was haemorrhaging cash and could not fill a Tuesday night.

The venues – and there were several, each more achingly beautiful than the last – shared one fatal characteristic. They had been designed from the outside in. The aesthetic was the strategy. Every operational decision had been subordinated to the look: kitchens compressed to give more floor space to the dining theatre, bar counters positioned for maximum visual impact but minimum restocking efficiency, bespoke design features that required specialist cleaning crews the profit margins could not support. Behind the scenes, the operations teams were fighting a daily war against their own venues – beautiful spaces that actively resisted the work of running a restaurant.

And the crowds? They came. Of course they came. The influencer circuit in any major city operates with the ruthless efficiency of a locust swarm: arrive, photograph, post, leave.[1] The venues were so overwhelmed by this initial wave of photogenic visitors that they never paused to build the one thing that keeps a restaurant alive past its first anniversary – a community of regulars who come on a wet Wednesday because they like the food and they like being there. When the Instagram set moved on to the next shining object, there was nothing underneath. The venues had themes. They had aesthetics. They had followers. What they did not have was a concept.

This distinction – between a theme and a concept – is where most hospitality ventures go wrong, and it is where this book begins.

Ideas Are Not Concepts

There is a moment, and it tends to strike between the second glass of wine and the bill arriving, when a perfectly intelligent human being looks around a busy restaurant and thinks: *I could do this*. The food is unremarkable. The service is clumsy. The décor hasn't been updated since the previous tenant gave up. And yet the place is packed. How hard can it be?

The answer, naturally, is that it can be extraordinarily hard – but not for the reasons most people assume. The restaurant industry does not fail operators because the work is physically demanding, although it is. It does not fail them because the hours are punishing, although they are. It fails them because the vast majority of new venues open with an idea rather than a concept, and the difference between those two things is the difference between a fantasy and a business.

An idea is a feeling. 'I want to open a wine bar.' 'I've always dreamed of running a beachfront restaurant.' 'There's nowhere good for brunch in my neighbourhood.' These are observations – sometimes accurate observations – but they are not commercial propositions. They describe what a venue looks like from the customer's chair. They say nothing about who sits in that chair, how often, at what price point, during which hours, and whether the economics of serving them can sustain a lease, a team, a supply chain, and a founder's mortgage.

A concept is the complete commercial engine underneath the aesthetic. It answers five questions simultaneously: who is the customer, what occasion does this serve, what is the price architecture, what is the experience promise, and what is the operational model that delivers all of the above at a sustainable margin? Remove any one of those five components and you do not have a concept. You have a mood board with a liquor licence.

Professor H.G. Parsa's longitudinal research across the American restaurant industry found that 26.16% of independent restaurants fail within their first year, with approximately 60% closing within three years. The single most frequently cited contributor was not undercapitalisation, not poor location, and not bad food. It was an ill-defined concept – operators who could not articulate, with commercial precision, what their venue was and who it was for. The Citrin Cooperman restaurant industry survey reinforced this from the

profit side: restaurants with clearly defined concepts and disciplined cost control achieved average earnings before interest, taxes, depreciation, and amortization (EBITDA) margins of 22%. Those in the middle managed 14%. Those at the bottom – the ones running on instinct and optimism – averaged 4%. The gap between a defined concept and a vague idea is not philosophical. It is eighteen percentage points of operating profit.

And here is the data point that should keep every aspiring operator honest: fine dining, the format that demands the most exacting concept definition, has the highest failure rate at 4.9% annually. Fast-casual, which typically operates within the tightest conceptual constraints – limited menus, standardized operations, obsessive unit economics – fails at 0.5–0.6%. The formats with the most romantic appeal are the ones that kill the most businesses. The formats that feel most constrained are the ones that survive.

Where Venue Concepts Come From

Viable hospitality concepts do not arrive in moments of inspiration. They emerge from one of four sources, and understanding which source yours comes from determines how you develop it.

The first is the **market gap** – a genuine, demonstrable absence in a specific location. Not 'there's no good Thai food in London', which is a statement so broad it means nothing, but 'there is no mid-priced Thai restaurant within the SE15 postcode that serves a lunch crowd of 2,000 office workers who currently have three sandwich chains and a Pret'. A market gap is defined by geography, price point, occasion, and customer density. If you cannot quantify all four, you do not have a gap. You have a hunch.

The second is **personal expertise** – a skillset or knowledge base deep enough to create genuine competitive advantage. The former head pastry chef at a Michelin-starred restaurant who opens a neighbourhood pâtisserie is not following a dream. She is deploying a specific, verifiable capability that her competitors in that postcode cannot match. Expertise-led concepts carry inherent credibility, but they carry a trap too: the assumption that excellence in craft automatically translates to excellence in commerce. It does not. The best cook in the building and the best operator in the building are frequently two

different people, and the concept must account for both.

The third is the **format transplant** – taking a proven model from one market and adapting it for another. Pizza Hut's expansion across the Middle East and Asia transformed the brand from an American quick-service restaurant (QSR) into a full-service family dining experience, because the cultural expectation of eating pizza in Riyadh or Kuala Lumpur was fundamentally different from grabbing a slice in Ohio. The concept was the same in name only; the format, service model, and price architecture were rebuilt from the ground up. Format transplants work when the operator genuinely understands both the origin market and the destination. They fail – spectacularly, reliably, and expensively – when someone eats a meal on holiday and thinks it will work back home without modification.[2]

The fourth is the **cultural moment** – capturing a shift in consumer behaviour before the market saturates. Khairul Aming, a Malaysian food content creator, opened Rembayung in 2026 as a RM4 million physical restaurant that generated four million digital engagements before serving its first plate. The concept worked because it was built on something the research literature calls trust-led demand: four years of content creation had already validated the audience, the price tolerance, and the emotional connection. The cultural moment was not the restaurant. The cultural moment was the relationship. The restaurant was simply its physical expression. But cultural moments are, by definition, temporary. Sweetgreen understood this: their seasonal menu strategy operates like a fashion drop, refreshing the concept at regular intervals to maintain the sense that something is always happening. Venues that capture a cultural moment without building a mechanism for renewal become Karen's Diner – a gimmick-based concept in the UK that collapsed under £400,000 of debt when the novelty expired. The cultural moment gave them an opening. The absence of a concept underneath it gave them a closing date.

The Concept Stress Test

Before you invest another hour – let alone another pound – your concept must survive five questions. Not the gentle, encouraging questions your friends will ask when you tell them about your exciting new venture over dinner. These questions:

Who is the customer? Not 'everyone' and not 'people who like good food'. A specific demographic with a specific income, a specific set of habits, and a specific reason to choose your venue over the fourteen alternatives within walking distance. The K-shaped economic recovery has split consumer spending into two diverging trajectories: the top 10% of households increased their discretionary spending by 62% over five years, while lower and middle-income spending stagnated or contracted. Your concept must know which side of that K it serves, because the operational model for each is entirely different. A venue that tries to serve both ends up in the squeezed middle – too expensive for the value-conscious, too ordinary for the affluent. Jamie's Italian learned this at a cost that eventually sank the entire chain.

What occasion does this serve? Lunch is not dinner. A birthday is not a Tuesday. A work meeting is not a date. The same person who will happily spend £85 per head on an anniversary will balk at £22 for a Thursday lunch. Your concept must identify its primary occasion and design everything – menu, pricing, service tempo, lighting, music – around the expectations of that specific moment.

Why here? A concept is not abstract. It exists in a specific physical location, within a specific competitive field, serving a specific catchment. 'This would work anywhere' is not a strategy. It is the absence of one.

Why now? Markets move. Consumer habits shift. UK dinner reservation averages have drifted to 6:12 PM; Hong Kong's peak dining time has migrated from 8 PM to 6 PM. If your concept is built around assumptions about when people eat, drink, and socialise that were true five years ago, you are opening a venue for a market that no longer exists.

What stops someone copying this in six months? If the answer is 'nothing' then you do not have a concept. You have a first-mover advantage with a six-month shelf life. The defensive moat for an independent venue is almost never the food or the décor – it is the depth of the operational model, the strength of the community, and the consistency of the experience delivered at speed under pressure. These are the things that take years to build and cannot be replicated by someone with a fit-out budget and a Pinterest board.

Format, Daypart, and Occasion

Here is something that almost no first-time operator grasps until it is too late: your venue is not one business. It is three, four, possibly five different businesses sharing the same address, the same kitchen, and the same electricity bill.

A coffee shop that opens at 6:30 AM and closes at 4:00 PM is earning up to 60% of its daily revenue in a three-hour morning window. The concept for that window – speed, consistency, minimal friction, grab-and-go packaging – is fundamentally different from the concept for the mid-morning lull, when the same space needs to become a laptop-friendly workspace for freelancers who will nurse a single flat white for two hours. Both are valid revenue streams. Both require deliberate concept design. Ignoring either means the venue is working at full capacity for three hours and burning cash for the remaining seven.

Clayton Christensen's 'Jobs to Be Done' framework, developed at Harvard Business School, provides the most useful lens for hospitality concept development. The now-famous milkshake study found that the same product was being hired for completely different jobs depending on the time of day: morning commuters bought milkshakes as a one-handed, slow-consuming breakfast substitute for a boring drive, while afternoon parents bought the same milkshake as a guilt-free treat for their children. Same product. Different job. Different customer. Different competitive set.[3]

Your venue operates on exactly this principle. The Friday-night couple celebrating an anniversary is hiring your restaurant for an entirely different job than the Wednesday-afternoon sales team booking a client lunch. The concept must account for every daypart you intend to trade, with a clear understanding of who the customer is, what they need, how long they will stay, and what they will spend. Cracker Barrel built a fifty-year American institution on this understanding, designing their format to serve the morning road-trip family, the weekday lunch retiree, and the Sunday-after-church crowd with equal conviction – same building, different jobs, differentiated menu architecture for each.

The danger of ignoring daypart strategy is not just lost revenue. It is concept incoherence. A venue that tries to be a buzzing cocktail bar at 10 PM and a serious business lunch destination at noon, without

THE SHIFT UP: OPEN

adjusting lighting, music, menu, service tempo, and pricing between those two identities, is not versatile. It is confused. And confused venues train their customers to be confused too – which means those customers default to somewhere simpler, where they know exactly what they are going to get.

THE DEBRIEF

Marcus Treamer - Hong Kong

The Error: I worked with a multi-concept operator in Hong Kong whose entire philosophy was design-first. The venues were undeniably beautiful – award-winning interiors, Instagram catnip, magazine features before a single plate had been served. But every operational decision was subordinated to the aesthetic. Kitchens were compressed. Bars were positioned for visual impact rather than workflow. Bespoke features required specialist maintenance the margins could not absorb. The concept was the design, and the design was the concept. There was nothing else underneath.

The Fallout: The influencer crowd arrived immediately and in force. The venues were slammed – which felt like success but was actually the worst possible outcome, because it meant the team never had breathing room to build regular clientele. They were too busy managing the Instagram rush to notice they had no Tuesday-night community forming. When the influencer cycle moved on – and it always moves on – the footfall vanished almost overnight. Beautiful venues, haemorrhaging cash, with no loyal customer base to sustain them.

The Fix: The lesson is structural. Design is not concept. Design serves concept. A venue that wins a design award and cannot fill a wet Wednesday has its priorities inverted. The concept – who the customer is, what occasion you serve, what keeps them coming back when the novelty fades – must be locked in before a single interior specification is written. Start with the commercial engine. Let the design express it. Never the other way around.

Killing Your Darlings

Every concept begins with enthusiasm. Someone has an idea they love – the rooftop terrace, the open-fire kitchen, the natural wine programme that will educate and convert an entire neighbourhood. And that enthusiasm is not worthless. It is, in fact, essential. Nobody survives the eighteen months between 'I have an idea' and 'we are open for business' without genuine conviction.

But conviction untempered by evidence is just optimism wearing a better outfit. And optimism, as this book will argue repeatedly, is the most expensive emotion in hospitality.

The cognitive biases working against you at this stage are well documented and relentless. Sunk cost fallacy will make you cling to elements of the concept you have already invested time and emotional energy in, regardless of whether they survive scrutiny. Optimism bias will lead you to overestimate demand and underestimate costs – research consistently shows that over 80% of startup founders miss their initial market share targets, and hospitality is no exception. Commitment bias will make you interpret ambiguous feedback as validation, because by the time you are deep enough into the process to receive feedback, you have already decided what you want to hear. These are not character flaws. They are structural features of the human brain, and they have sunk more restaurants than bad food ever has.

The discipline of killing your darlings is the discipline of separating what excites you from what the market will sustain. It does not mean abandoning creativity. It means subjecting creativity to the same scrutiny you would apply to any other commercial decision. The rooftop terrace is spectacular – but does the lease allow external trading past 10 PM, and what does the revenue model look like when it rains for four months? The open-fire kitchen is romantic – but does the extraction system meet local compliance, and have you costed the fuel, the insurance premium increase, and the additional cleaning hours? The natural wine programme is admirable – but does your target customer understand it, want it, and will they pay the premium it requires to achieve an acceptable pour cost?

This is not pessimism. It is due diligence with a personality. The operator who can interrogate their own concept with genuine rigour –

who can sit with the discomfort of discovering that their favourite element is commercially unviable – is the operator who opens a venue that lasts. The one who cannot is the one who opens a venue that looks magnificent in the photographs taken during the first month and appears on a commercial lease assignment listing eight months later.

The Concept Canvas at the end of this chapter will force these confrontations. It is a single page, and everything on it must fit. If your concept cannot be expressed clearly in that space – concept name, target customer, primary occasion, price positioning, experience promise, competitive differentiation, and a three-sentence elevator pitch – then it is not yet a concept. It is still an idea. And ideas, however beautiful, do not pay rent.

A completed Concept Canvas does not mean your concept is ready. It means your concept exists – as a defined, articulated commercial proposition rather than an atmospheric fantasy. It can now be tested, challenged, costed, and improved. It can be handed to a business partner, an investor, or a designer, and they will understand what you are building and why.

That testing begins in the next chapter, where the concept meets the market for the first time – not with a lease or a fit-out, but with evidence. Because the most dangerous thing about a beautifully defined concept is the temptation to believe in it before anyone else has confirmed they would pay for it.

[1] The average lifespan of a venue as a 'trending' Instagram location in a major city is approximately 12–16 weeks. This is not an insult to influencer marketing – it is a structural characteristic of attention-based media. A venue designed to capture attention and a venue designed to retain customers are solving two different engineering problems, and they require two different concepts.

[2] The format transplant failure rate increases dramatically when the operator's experience of the origin market is limited to having been a customer there. Eating paella in Barcelona does not constitute market research for a Spanish restaurant in Birmingham. This should be obvious. It frequently is not.

[3] Christensen's milkshake study remains the single most useful framework for hospitality concept development, despite having nothing

THE SHIFT UP: OPEN

to do with hospitality. The original research was commissioned by McDonald's, who wanted to know how to sell more milkshakes. The answer – stop thinking about milkshakes and start thinking about what job the customer is hiring the milkshake to do – applies to every daypart decision you will ever make.

CHAPTER 2

Know Your Market – Research That Prevents Expensive Assumptions

THE SHIFT UP: OPEN

In Shanghai, in a mixed-use development designed by the architect John Portman, a cocktail bar occupied a space that had previously housed the Long Bar – an institution so embedded in the city's expatriate drinking culture that its closure made the English-language papers. The new venue, CVRVE, had a cocktail programme led by Cross Yu, who would go on to win the Diageo World Class competition. The Portman Ritz-Carlton centre contained a five-star hotel, luxury residences, office towers, and an established retail and food and beverage (F&B) mix. On paper, every indicator pointed to an A-grade location with guaranteed multi-daypart footfall.

There was one escalator between the ground floor and the venue.

No demographic report would have flagged it. No footfall study would have quantified its effect. No spreadsheet capturing the 14,000 office workers in the adjacent towers would have noted that precisely zero of them would take an escalator upward to discover a bar they did not already know existed. The evening crowd – people who came specifically for CVRVE, who had heard about the programme and sought it out – filled the room nightly. The lunch daypart, which the financial projections had assumed would capture office workers looking for a quick, quality meal in a premium setting, never materialised. Within ninety days, the venue had pivoted to opening daily at 5 PM. Not because the concept was wrong, but because the market – the real, physical, walkable, observable market – was telling a different story than the one the data had promised.

This is what market research actually means. Not the data. The data is the beginning. Market research is what happens when the data meets the pavement.

What Market Research Actually Means

Most aspiring operators treat market research as a task to complete – a box to tick before writing the business plan. They spend an evening with Google, gather some demographic data, note that there are 35,000 residents within a two-mile radius, observe that the area 'seems to be growing', and declare the market validated. This is not research. This is confirmation bias with a browser history.

Genuine market research answers seven questions, and it answers them with evidence rather than instinct. Who lives here, works here, and visits here? What do they currently spend on eating and drinking out, and where? What is missing from the current offering – and is that absence an opportunity or a warning? What price points does this market support? What is the competitive density, and where are the genuine gaps? How is this market moving – growing, contracting, gentrifying, saturating? And what are the physical, logistical, and regulatory realities that no data source will tell you?

The consequence of skipping this process is not abstract. CB Insights' post-mortem analysis of failed startups across all sectors found that 42% of businesses that closed cited 'no market need' as the primary cause.[1] Not insufficient capital. Not poor execution. The fundamental absence of a customer base willing to pay for what was being offered. In hospitality, where margins are measured in single digits and the lease does not care whether anyone turns up, opening without market evidence is not optimistic. It is reckless.

And the definition of 'market' has changed. Seventy per cent of online searches now conclude without the user clicking through to any website – the so-called zero-click search. Your potential customer is making decisions about where to eat, drink, and spend money based on the information that appears in a search results page, a Google Maps listing, or an AI-generated summary. If your market research only considers the physical catchment – who walks past, who lives nearby – you are studying half the market. The other half is digital, and they are making choices before they leave the house. Understanding how your target customer discovers venues in your market is now as critical as understanding whether they exist at all.

Desk Research: Demographics, Demand, and Data

Desk research is the foundation – necessary but not sufficient. It establishes the quantitative parameters of your market before you set foot in the area, and it prevents the most expensive category of error: opening in a location that cannot mathematically support your concept.

The critical calculation is revenue per square foot. Full-service restaurants in developed markets require a minimum of approximately \$150 per square foot annually to cover occupancy costs, labour, food,

and a survival-level margin. Fast-casual models, with their higher throughput and lower labour ratios, need closer to \$500 per square foot because the average transaction value is lower and the model depends on volume. If the available premises, at the rent being quoted, cannot generate these minimums based on conservative seat-turn and average-spend assumptions, the location fails before you order the first case of wine.

The data sources vary by market, and knowing where to look is half the battle. In the UK, the Office for National Statistics (ONS) provides local-area demographic profiles, and local authority planning departments publish footfall studies, tourism statistics, and commercial vacancy rates – most of it free, most of it ignored. In the US, the Census Bureau's American Community Survey and the Bureau of Labor Statistics' Consumer Expenditure Survey provide household spending data at the county and metropolitan level. In the UAE, the Dubai Department of Economy and Tourism publishes restaurant density and tourism arrival data quarterly. In Thailand, the Tourism Authority and the Board of Investment provide regional visitor statistics, though the data requires careful interpretation – a province-level tourist count tells you very little about whether those tourists eat in your postcode.

For deeper market intelligence, Euromonitor Passport reports cost approximately \$1,475 per market report. Technomic Ignite provides US-specific restaurant industry data at institutional-research prices. Statista offers a mid-tier alternative. These are not optional extras for operators planning to invest six figures. They are the equivalent of a structural survey before buying a building – except the building is made of perishable food and staffed by humans who expect to be paid every fortnight.

The global hospitality market reached approximately \$4.9 trillion in 2024. The fact that the market is enormous does not mean your corner of it will sustain your venue. National growth figures are irrelevant at the postcode level. A city growing at 3% per annum may contain neighbourhoods contracting at 5%. The menu price inflation across the US restaurant industry in late 2024 averaged just 1.41% – which means operators who built their projections on the assumption that they could raise prices to offset rising costs were working with margins thinner than a crêpe and twice as fragile.

One further data point deserves attention before you leave the desk for the street. Franchise operations targeting underserved niche markets – neighbourhoods where a specific cuisine, price point, or format is demonstrably absent – report 35% higher success rates than those entering established competitive markets. The lesson for independents is identical: the granularity of your market data determines the quality of your positioning. 'There are lots of restaurants in this area' is not an analysis. 'There are fourteen restaurants within 500 metres, eleven of which serve food priced between £12 and £18 per head at lunch, and none of which offer a dedicated takeaway-and-delivery model for the 3,200 office workers who leave the building between 12:15 and 12:45' is an analysis. The difference between those two statements is the difference between a hunch and a strategy.

Field Research: Feet on the Ground

This is where most operators either distinguish themselves or doom themselves, and the method is so simple it feels almost insulting to state: go there. Stand where your customers would stand. Walk the route they would walk. Do this on a Tuesday lunchtime, a Friday evening, a Sunday morning, and a rainy Wednesday afternoon. Do it three weekends running.

Record everything. Photograph competitor queues at peak times, note which side of the street carries the heaviest foot traffic, log the precise time the lunch rush begins and ends. Anecdotal impressions fade within days; documented observations compound into evidence.

The Three Weekends methodology is not a published academic framework. It is a practitioner discipline – the minimum viable observation period that accounts for weekly trading variation, weather effects, and the difference between what a street looks like when you visit it once and what it actually does across a representative sample. During those three weekends, you are counting footfall at your prospective site (not the street – the site), timing pedestrian dwell, observing the trading patterns of your nearest competitors, and noting every physical feature that would affect customer access: stairs, escalators, parking availability, signage visibility from the main thoroughfare, the direction people naturally walk.

THE SHIFT UP: OPEN

The escalator at the Portman Ritz-Carlton killed an entire daypart. A similar physical barrier – a set of steps, a poorly positioned entrance, a car park that faces away from the high street – will do the same to yours. No desk research will reveal this. Only feet on the ground.

Customer intercept research formalizes what observation begins. Two hundred structured responses – brief, specific, conducted at different times across different days – provide a 7% margin of error at 95% confidence. That is sufficient to identify whether your target customer exists in this location, what they currently spend, what they feel is missing, and what would bring them through your door. The questions must be direct and behavioural, not aspirational: 'How often do you eat out within a ten-minute walk of here?' produces useful data; 'Would you visit a new Mediterranean restaurant in this area?' produces polite fiction. People will tell you what they do. They will lie – with the best of intentions – about what they might do.[2]

Competitive auditing during field research goes beyond counting restaurants. For every direct competitor within your defined catchment, note their price points (photograph the menu), estimate their occupancy at three different times across three different days, observe their service model and speed, and identify their apparent target customer. For indirect competitors – the supermarket meal deal, the office canteen, the Deliveroo dark kitchen – note their pricing and convenience advantage. Your concept does not compete only with other restaurants. It competes with every alternative your target customer has for solving the same job at the same time of day.

Pay particular attention to what the competitors are doing badly. An empty restaurant at 7 PM on a Friday is a data point. A restaurant with a thirty-minute wait at 12:30 on a Wednesday is a data point. A venue with strong reviews but visible operational chaos – long wait times, confused service, inconsistent food – is a data point that tells you this market has demand that is currently being served poorly. That is the most valuable competitive gap of all, because it means customers are already trained to eat out in this location. They just need a reason to choose you instead.