Amazon Business Growth Machine: Intent-Centric Edition

The Complete Blueprint for Modern Amazon Selling

HOW TO LEVERAGE INTENT ANALYTICS, BRAND ANALYTICS, AND AI TO SYSTEMATICALLY OUTPERFORM THE MARKET

> *For every Amazon seller who optimized individual keywords and never understood why a competitor was quietly capturing entire customer intents.*

AMAZON BUSINESS GROWTH MACHINE: INTENT-CENTRIC EDITION

The Complete Blueprint for Modern Amazon Selling

HOW TO LEVERAGE INTENT ANALYTICS, BRAND ANALYTICS, AND AI TO SYSTEMATICALLY OUTPERFORM THE MARKET

For every Amazon seller who optimized individual keywords and never understood why a competitor was quietly capturing entire customer intents.

Preface

This is not a revision of the original Growth Machine book. It is an evolution.

The first edition taught sellers to think about keywords. This edition teaches you to think about *customer intent*.

The difference is not semantic. It is structural.

A keyword is a string of words: "memory foam pillow," "orthopedic pillow for neck pain," "pillow for side sleepers." These are tactical. They are what a shopper types into the search bar.

An intent is the underlying goal: "I need a pillow that reduces neck pain." Thirty different shoppers might express this intent using 30 different keywords. Traditional analytics force you to analyze each keyword in isolation. You optimize the title for "orthopedic pillow." You bid on "pillow for neck." You create A+ content for "cervical support." Three separate strategies for one intent.

Intent-centric analytics do something different. They cluster the 30 keywords into a single intent bucket. They aggregate the performance metrics across all 30 terms. They reveal the true size of

the market you are competing for. And they point you toward a unified strategy that wins on the entire intent, not individual keywords.

This shift — from term-centric to intent-centric — is the next evolution in Amazon selling. Amazon itself is moving here with Alexa for Shoppers and intent-based search. The sellers who make this shift first will have a two-year competitive advantage before the rest of the market catches up.

This book shows you how.

Who This Book Is For

- **Brand-registered Amazon sellers** who have used keyword analytics before and are ready for a higher level of sophistication.
- **Amazon advertising managers** who want to move beyond ACoS optimization to intent-level market-share strategy.
- **E-commerce operators** who need to connect intent-driven marketing decisions to unit economics and inventory constraints.
- **Technical sellers** who are ready to use AI-powered agents and NeonPanel's intent tools to compress analysis time from hours to minutes.

You do not need to be a programmer. But you need to be willing to think differently about your market.

How to Read This Book

This book is structured in seven sections, building on the first edition but with intent-centric frameworks throughout.

Section I: The Data Engine introduces intent-centric thinking. What intent is, why it matters, how it differs from keyword analysis, and how NeonPanel's intent tools work.

Section II: Intent-Driven Marketing & Advertising is the analytical core. Four chapters on finding intent gaps, measuring advertising incrementality, discovering blue-ocean intents, and building your intent strategy.

Section III: Operations & Profitability grounds intent analysis in financial reality. Your unit economics and inventory constraints determine *how aggressively* you can pursue intent opportunities.

Section IV: The Weekly Governance Loop turns intent analysis into a perpetual system. The ritual, the cadence, the feedback loop.

Section V: The Persistent Brain covers long-term memory — how your machine learns and improves across weeks and months.

Section VI: The NeonaSphera Agent provides the AI orchestration specification for automating the intent-centric Growth Machine.

Section VII: Practical Workshop is hands-on. You run the machine on your own product, using real data and real intent tools.

Each chapter can be read independently, but the power is in the system. Read straight through on the first pass, then return to individual chapters as reference.

INTRODUCTION: THE INTENT REVOLUTION

Why This Book Exists, and Why the Timing Matters

You have read every Amazon seller guide. You have optimized your Advertising Cost of Sale (ACoS) until it is pristine. Your ad console is a masterpiece of efficiency. And yet something feels incomplete.

You are watching revenue plateau while a competitor you have barely heard of is quietly dominating your market. Not because they have lower prices or better products, but because they understand something you don't: the market doesn't think in keywords. It thinks in *intent*.

The Keyword-Centric Trap

For fifteen years, Amazon selling was governed by a simple framework: find high-volume keywords, optimize for rank, bid aggressively, measure ACoS.

This framework was never wrong, exactly. It just became incomplete.

When Amazon's algorithm was simpler — when rank position directly corresponded to sales — keyword-level optimization worked. You bid on "memory foam pillow," tracked your rank position, and celebrated when it hit page one. Done.

But Amazon's algorithm evolved. Twice, in fact.

First came Category Relevance. Then came the A10 algorithm, which stopped looking at individual keywords and started looking at *topical clusters* — groups of related keywords that indicate comprehensive market understanding. The algorithm learned to ask: "Does this seller understand the full landscape of customer needs in this category? Or are they optimizing for one keyword at a time?"

When you optimize for one keyword, you appear relevant to one keyword.

When you optimize for an *intent* — when you understand that "memory foam pillow," "orthopedic pillow," "pillow for neck pain," and "cervical support pillow" all express the same underlying customer goal — you appear authoritative across an entire market segment. Amazon's algorithm notices. It rewards that with visibility, organic rank, and conversion opportunity.

Helium 10, Perpetua, Sellersprite — the tools built over the last decade — are fundamentally keyword-centric. They track rank position on individual keywords. They tell you which keyword is "winning." But they never step back and ask: *What is the market trying to tell you across all 30 variations of the same customer need?*

That blind spot is the opportunity.

The Market Has Already Shifted

You are not the only one who feels the incompleteness. The entire industry is moving.

67% of B2B marketers now use intent data for advertising decisions (2024 survey). That number will be higher by the time you read this.

The intent-based marketing technology market is growing at 839% per decade — from $2.9 billion in 2025 to $27.3 billion by 2035. McKinsey found that intent-driven personalization delivers **5–8x the ROI of generic campaigns**.

Amazon itself is moving toward intent. Alexa for Shoppers, Amazon's generative AI shopping assistant, does not think in keywords. You ask it a natural question ("I need a pillow that doesn't retain heat and supports my neck"), and Alexa for Shoppers interprets your *intent* and surfaces relevant products. It sees "cooling," "neck support," and "sleep comfort" as semantic clusters, not as three separate keyword opportunities.

Your competitors are watching this shift. Some are already moving. The sellers who understand intent clustering will have a 24–36 month advantage before the rest of the market catches up.

By the time everyone else figures it out, you will have built a compounded knowledge advantage that they cannot replicate.

What Academic Research Says

This is not speculation or marketing trend-chasing. Peer-reviewed research validates every core principle in this book.

On Intent Prediction:

Machine learning models can predict customer purchase intent with **92.39% accuracy** using behavioral signals (session duration, page depth, interaction patterns, search history). This is not experimental — it is proven and deployed at scale by Shopify, Wayfair, and Amazon itself.

Source: IEEE Conference on Customer Purchase Intent Prediction using Feature Aggregation, 2024; ResearchGate studies on ML-driven intent modeling, 2025

On Segmentation:

Natural Language Processing applied to customer behavior produces what researchers call "intent-driven segmentation." Companies implementing NLP-based intent clustering report:

- **30–50% improvement in email engagement metrics**
- **10%+ lift in sales**
- **Measurable improvements in customer loyalty and lifetime value**

Case study: Sur La Table implemented AI-powered intent-aware search optimization and saw:

- **6.6% increase in add-to-cart rates**
- **11.5% improvement in category average order value**

Source: ScienceDirect, Springer Nature peer-reviewed research on NLP in marketing; CleverTap research citing McKinsey and real-world implementations, 2024

On Semantic Clustering:

Grouping keywords by semantic meaning and SERP intent (not by exact match or volume alone) expands your addressable market. Sellers who organize campaigns and content around intent clusters rather than individual keywords capture wider audience share.

Amazon's A10 algorithm explicitly interprets content as "topical clusters indicating comprehensive market coverage." When your listing demonstrates authority across related customer needs, the algorithm interprets this as higher conversion probability and rewards visibility accordingly.

Source: Eva.guru, Amazon Keyword Clustering Guide 2025; SeRanking Comprehensive Guide to Keyword Clustering; InfraNodus Keyword Clustering with AI

The Competitive Moment

Here is what is happening right now, in May 2026:

Market Leaders Are Waking Up:

- Amazon is pushing Alexa for Shoppers (intent-driven search) aggressively.
- Advertising tools (Perpetua, Bidx, Helium 10) are starting to add "intent" features, but retrofitted onto keyword-centric architecture. Half measures.
- Early-mover sellers and agencies are building internal intent systems.

The Window is Open:

- Most sellers still think in keywords.
- Most tools are still keyword-centric.
- Intent-native platforms barely exist.
- You have 18–24 months before this becomes table-stakes.

After That:

- Expect Helium 10 to fully integrate intent (and charge more for it).
- Expect Amazon to surface Alexa for Shoppers intent signals in Seller Central.
- Expect every new seller to be taught "think in intents, not keywords."
- Expect your competitors to catch up.

The Question: Do you want to be ahead of the wave, or riding it after everyone else?

What This Book Teaches

This book is not a tutorial on using a specific tool (though it heavily features NeonPanel, a purpose-built intent-native platform). It is a **strategic framework for modern Amazon selling**, grounded in customer psychology, machine learning research, and proven business outcomes.

Section I: The Data Engine

You will learn what intent is, why it matters, and how NeonPanel's intent clustering (both LLM-assisted and manual) transforms raw keyword data into strategic insight. You will run your first intent analysis and see your business through a new lens.

Section II: Intent-Driven Marketing & Advertising

You will learn the four frameworks that drive growth:

- **Intent Gap Analysis:** Finding high-volume intents where your brand share is below 10% (your biggest opportunities)
- **Incrementality at Intent Level:** Proving whether your advertising is truly incremental or cannibalizing organic sales (prevents wasted spend)
- **Blue Ocean Intent Discovery:** Identifying under-competed intents with high ROI (the profit machine)
- **Intent-Based Budget Allocation:** Moving money from low-opportunity to high-opportunity intents (the efficiency lever)

Section III: Operations & Profitability

You will ground intent strategy in financial reality. True unit economics, inventory constraints, and the "marketing throttle" that prevents you from destroying profitability while chasing growth.

Section IV: The Weekly Governance Loop

You will learn the ritual that turns intent analysis into a compounding system. Every Monday, you run the same sequence. Every week, you measure outcomes, refine hypotheses, and learn. By week 12, your machine is not guessing — it is predicting.

Section V: The Persistent Brain

You will build long-term memory. Action logs. Hypothesis tracking. Over 24 weeks, you will accumulate proprietary knowledge about your specific market that no competitor can replicate because they are stuck in keyword-level analysis.

Section VI: The NeonaSphera Agent

You will see how to automate the analytical layer. An AI agent pulls data, applies frameworks, checks constraints, and generates recommendations — leaving strategic decisions to you. Analysis that used to take hours now takes 15 minutes.

Section VII: Practical Workshop

You will run your first complete intent analysis on your own product, using real NeonPanel tools and real data. By the end of Chapter 13, you will have a complete "intent dossier" for your business.

Chapter 14: Case Study

You will see a real seller (pillow brand, mid-market) who was plateaued with keyword-centric optimization. They switched to intent-centric thinking. In 8 weeks, revenue grew 36%. Their cost

per acquisition dropped. Their market share on key intents nearly doubled. This is not theoretical. This is what happens when you think in intent.

Who Should Read This Book

You should read this if:

✓ You are brand-registered on Amazon with at least 3 ASINs in the same category.

✓ You are spending >$5K/month on advertising and wondering if you are leaving money on the table.

✓ You have optimized individual keywords to death and the improvements have plateaued.

✓ You have access to Brand Analytics but feel like you are only seeing the trees, not the forest.

✓ You want to move beyond "improve ACoS" to "own this market."

✓ You are comfortable with data and willing to change how you think about your business.

You might not need this if:

✗ You are selling on Amazon as a side project (<$20K/month revenue).

✗ Your product is so unique that you have no real competition.

✗ You have zero interest in using AI or learning new tools.

A Note on Tools

This book features NeonPanel extensively. NeonPanel is a data integration platform that connects to your Amazon Seller Central account and exposes Brand Analytics, advertising data, financial data, and inventory data through AI-friendly tools.

Here is the honest truth: The intent-centric *frameworks* in this book can be executed with any data source. You can manually cluster keywords using a spreadsheet, pull SQP data monthly, and run the analyses without any special tool.

But speed matters. A seller who runs the weekly governance loop on Monday morning using NeonPanel's intent tools (15 minutes of analysis) will outcompete a seller who runs it quarterly using manual spreadsheets (8 hours of work). The first seller learns 52 times per year. The second learns 4 times per year.

The data is the same. The frameworks are the same. The execution velocity determines the outcome.

If you already have a tool stack that works, some of these frameworks will translate. If you want to move fast and systematically, NeonPanel is purpose-built for this exact workflow.

What You Will Gain

By the end of this book, you will:

1. **Understand how your market actually works.** Not "this keyword is popular" but "this intent represents 247K monthly searches, I own 7% of it, and I am leaking share because my conversion rate is weak."

1. **Identify your biggest leverage points.** Not "bid more on profitable keywords" but "the Pillow for Hot Sleepers intent is a hidden gem with 3.2x ROAS; I am underinvesting; if I shift $4K/month there, I will see $12K/month uplift."

1. **Stop wasting money.** Not "this keyword has high ACoS" but "my brand is so strong organically on this intent that advertising is cannibalistic; reduce spend by 30% and keep the sales."

1. **Build a learning system.** Not "run analysis once and implement" but "run the loop every week, measure outcomes, refine the model, compound improvements."

1. **Own your competitive position.** Not "rank higher" but "understand the playing field, where I am strong, where my competitors are vulnerable, and how to build defensible advantage."

1. **Free your brain for strategy.** Automation handles the data pulling and recommendation generation. You focus on decisions: which intents to defend, which to attack, how to allocate capital.

The Competitive Advantage Window

The seller who adopts intent-centric thinking in June 2026 will have an unfair advantage over competitors stuck in keyword-centric thinking.

By 2028, when everyone else has adopted it, the advantage is gone. But by then, you will have compounded 24 months of intent learning. You will know your market better than anyone. You will have built a machine that learns and improves every week.

That compounded advantage is defensible.

How to Use This Book

First Read (Understanding):

Read straight through. Chapters build on each other. By the end of Section II, you will understand the intent-centric framework. By the end of Section VII, you will have executed it.

Reference Use (Implementation):

After the first read, return to individual chapters as needed. Chapters 3–6 are your "how to get started" reference. Chapters 9–10 are your "weekly ritual" reference. Chapter 13 is your "data extraction" reference.

Tool Guide:

The Integration Guide (companion document) provides detailed specifications for every NeonPanel intent tool, with inputs, outputs, and workflows. Pair it with the relevant book chapters.

Why Now

Amazon's market is becoming more competitive, not less. Your customers are more sophisticated. Your competitors are smarter. The days of incremental keyword-level optimization delivering double-digit growth are ending.

The sellers who will win in the next 24 months are the ones who shift from tactical keyword optimization to strategic intent ownership.

This book teaches you how.

Let's Begin

The first step is understanding the data that is already available to you — data most sellers ignore.

Turn the page. We start with Brand Analytics, the most powerful tool you are not using.

SECTION I: THE DATA ENGINE

Chapter 1: Beyond ACoS — The Intent Paradigm

WHY INTENT-BASED ANALYTICS IS THE NEW NORTH STAR FOR AMAZON SELLERS

For years, Amazon advertising was governed by a single metric: Advertising Cost of Sale. Sellers optimized bids, pruned keywords, and celebrated when ACoS dropped. Agencies built entire business models around it. "We got your ACoS from 25% down to 18%" was the headline of every case study.

The problem? ACoS tells you how efficiently you spent money. It tells you nothing about whether you are winning.

Consider this scenario. Your ACoS is a pristine 15%, and your campaigns are "profitable." Your agency is sending you green dashboards and congratulatory emails. Meanwhile, a competitor — one you have never even looked at — has quietly captured 40% of purchase share across all the keywords in your most important category. They own the entire *intent* — the cluster of customer needs that your product addresses. They are building brand recognition at industrial scale. They are accumulating reviews that compound their conversion rate. They are compressing your organic rank, position by position, week by week.

By the time your ACoS dashboard reflects the damage, the market has already shifted beneath your feet.

ACoS is a rearview mirror. It tells you what happened inside your campaigns. It says nothing about:

- How many people searched for your most important keyword last month.
- What percentage of those shoppers saw your product.
- What percentage clicked.
- What percentage bought from you versus from someone else.
- Whether your market share is growing, stagnant, or shrinking.
- And — the question that determines your future — *what underlying customer need unites all of those shoppers?*

Brand Analytics answers every one of these questions. Intent analytics answers the last one.

WHAT BRAND ANALYTICS ACTUALLY IS

Amazon Brand Analytics is a suite of reports available to brand-registered sellers through Seller Central. It draws on Amazon's first-party data — the searches, clicks, and purchases of hundreds of millions of shoppers — and exposes aggregated, anonymized views of market-level behavior.

The three reports that matter most:

Search Query Performance (SQP). For any search term relevant to your product, SQP shows total market activity — searches, impressions, clicks, cart adds, purchases — alongside your brand's share of each. This is the core of the Growth Machine.

Search Catalog Performance (SCP). Starting from your ASIN rather than a keyword, SCP shows your full conversion funnel — impressions to clicks to carts to purchases — compared to the market. This is your diagnostic tool for understanding *where* your funnel breaks.

Top Search Terms. A ranking of the most popular search terms in your category by search volume. This is your discovery tool for finding keywords — and eventually, intents — you did not know existed.

THE INTENT PARADIGM SHIFT

When you shift from "How cheaply can I get a click on this keyword?" to "What share of this customer intent do I own?", your entire decision-making framework changes.

Bid strategy becomes intent strategy. Instead of asking "What should I bid on 'memory foam pillow'?", you ask "What purchase share do I want on the entire 'Pillow for Pain Relief' intent — which encompasses 30 keyword variations and 247,000 monthly searches — and what resources are required to achieve it?"

Budget allocation becomes intent investment. Instead of distributing budget evenly across campaigns, you concentrate it on high-value intents where your share is low and the market is growing.

Keyword targeting becomes intent targeting. You stop mining individual keyword variants and start clustering related terms into coherent customer goals. You stop asking "which keyword?" and start asking "which customer need?"

"Profitability" gets redefined. A keyword with 22% ACoS that is growing your purchase share on a 500,000-search-per-month intent may be more valuable than a keyword with 8% ACoS on a 2,000-search-per-month intent. Context — intent size, growth, competitiveness — changes everything.

HOW INTENT ANALYTICS WORKS

An intent cluster is a group of related search terms that represent a single customer goal. Example:

Intent: "I need a pillow that reduces neck and shoulder pain."

Search terms that express this intent:

- "orthopedic neck pillow"
- "pillow for neck pain and shoulder pain"
- "cervical support pillow for sleeping"
- "contour pillow for neck alignment"
- "therapeutic neck pillow"
- And 20 or more additional variations.

Traditionally, you would analyze each term separately — 25 separate rows in a spreadsheet, 25 separate bid decisions. With NeonPanel's intent tools, you cluster all 25 terms into one intent, aggregate their Brand Analytics data, and see the true market picture:

- **Total monthly searches across all 25 terms:** 247,000
- **Your brand's total purchase share across all 25 terms:** 7%
- **Competitor A's purchase share:** 22%
- **Competitor B's purchase share:** 15%
- **Fragmented rest of market:** 56%

Now you know: you are competing for a 247K-search-per-month intent where you hold only 7%. This is a strategic priority. A unified campaign targeting the entire intent is more powerful than 25 separate micro-optimizations.

THE CORE THESIS

An Amazon business is a machine. It has inputs (inventory, capital, content), processes (advertising, organic ranking, conversion), and outputs (revenue, profit, market share). Like any machine, it can be tuned, optimized, and scaled — but only if you have the right instruments.

ACoS was the speedometer. It told you how fast you were going. Useful, but not sufficient.

Brand Analytics is the GPS. It tells you where you are, where the market is, and which roads lead to growth.

Intent analytics is the strategic overlay — the layer that groups those roads into territories, shows you which territories are worth defending, and reveals which ones your competitors are leaving unguarded.

This book is the engineering manual for building a machine that uses all three instruments.

THE THREE TIERS OF NEONPANEL INTENT TOOLS

NeonPanel provides ten intent-centric tools organized in three tiers:

Tier 1: Intent Lifecycle — Create intent clusters manually or via AI. Map search terms to intents. List and audit your intent library.

Tier 2: Intent Analytics — Analyze intent-level performance with aggregated SQP metrics. Compare intents side-by-side. Visualize intent and ASIN performance across your full product range. Rank intents by opportunity type.

Tier 3: Intent Playbooks — Prescribe intent-level actions tied to opportunity classification. Extend the growth machine's diagnostic logic to intent granularity.

By the end of this book, you will be fluent in all three tiers.

Chapter 2: The Golden Trio + Intent Clustering

FOUR DATA LAYERS THAT TOGETHER CREATE A COMPLETE PICTURE

The original Growth Machine ran on three data feeds: SQP, SCP, and PPC Search Terms. This edition adds a fourth — Intent Clusters — the organizational layer that transforms raw keyword data into strategic insight.

Think of it this way:

- **Search Query Performance (SQP)** is the market terrain. Every keyword, every metric, every competitor's position.
- **Search Catalog Performance (SCP)** is your funnel scan. Where are you losing customers between impression and purchase?
- **PPC Search Terms** is your financial ledger. What are you paying for, and what return are you getting?

- **Intent Clusters** are the map that organizes the terrain. They group raw keyword data into coherent customer territories and reveal which regions of the market matter most — and where you should concentrate.

When you overlay all four, you see the full picture: market structure, your competitive position within it, where your funnel leaks, the financial cost of your response, and — for the first time — the strategic shape of the battlefield.

2.1 SEARCH QUERY PERFORMANCE (SQP)

SQP answers the most fundamental question in Amazon business: for a given search term, what slice of the market does my brand own?

When a shopper types "memory foam pillow" into the Amazon search bar, a cascade of events follows. Amazon shows results. The shopper scans the page. They click on one or more products. They add some to their cart. They buy one. SQP captures all of this — for every product on the page, aggregated by brand.

Search Volume. The total number of times this search term was entered in the period. This is the size of the market on this keyword. A keyword with 100,000 monthly searches is a highway. A keyword with 500 searches is a country lane. Both can be valuable, but they require different strategies.

Total Impressions vs. Brand Impressions. An impression occurs when a product appears in the search results for this term. Total impressions tell you how many product views the keyword generated across all sellers. Brand impressions tell you how many of those were yours. The ratio is your Impression Share — a measure of raw visibility.

Total Clicks vs. Brand Clicks. Of all the shoppers who clicked on a product from these search results, how many clicked yours? Click share measures attractiveness — your main image, title, price, rating, and badge status all influence whether a shopper clicks your listing versus a competitor's.

Total Cart Adds vs. Brand Cart Adds. Cart share measures consideration. A customer who adds your product to their cart has moved past casual browsing and is actively evaluating. If your click share is high but your cart share is low, something on your product detail page is losing the comparison — price, reviews, delivery speed, or content quality.

Total Purchases vs. Brand Purchases. This is the moment of truth. Purchase share is the percentage of all orders on this keyword that belong to your brand. It is the single most important number in the Growth Machine, because it directly measures your ability to convert market demand into revenue.

In the intent-centric framework, you no longer ask how you perform on one keyword. You ask how you perform across the entire cluster of keywords expressing the same customer goal. That aggregated purchase share — across all 25 or 30 related terms — is your true competitive position on the intent.

2.2 SEARCH CATALOG PERFORMANCE (SCP)

If SQP starts with a keyword and asks "who wins?", SCP starts with your product and asks "where does it win and lose?"

SCP shows the full conversion funnel for your ASIN: Impressions to Clicks to Cart Adds to Purchases.

Each step is a potential leak. The power of SCP is not in the absolute numbers — it is in the ratios between stages, because each ratio diagnoses a specific type of problem.

High Impressions, Low Clicks. Your product is being shown, but shoppers are scrolling past it. The diagnosis is almost always a listing quality issue above the fold: main image, title, price, or rating. This is the one area where you compete before the customer ever visits your page. If your click-through rate is significantly below the category average, no amount of advertising spend will fix the root cause.

Practical fixes: test a new main image (lifestyle vs. product-on-white vs. infographic). Shorten and restructure your title for mobile readability. Check if your price is competitive. Verify your rating is not below the category threshold (typically 4.0 stars).

High Clicks, Low Cart Adds. Shoppers are visiting your listing but leaving without adding to cart. The diagnosis is comparison failure: something on your product detail page loses to alternatives the customer is evaluating in parallel. Common culprits: price higher than similar products without clear value differentiation, review count low relative to established competitors, bullet points and A+ content that do not address the customer's primary concern, or delivery dates slower than Prime expectation.

High Cart Adds, Low Purchases. This is the most frustrating leak because the customer was almost there. They liked your product enough to cart it, then chose something else at checkout. Often the cause is a competitor offering a last-minute incentive — a coupon, a Subscribe and Save discount, or a Lightning Deal. Sometimes it is delivery speed: a customer who carts two similar products will choose the one arriving tomorrow.

The ratio between cart adds and purchases is your close rate. Track it weekly. If it drops, something in the competitive landscape changed.

In intent context, SCP becomes even more powerful. When you see a click-to-cart leak across all your ASINs competing in the same intent, you know the problem is not a single listing — it is a

systematic gap in how you address that customer need. That signals a content strategy issue, not a bid strategy issue.

2.3 PPC SEARCH TERMS

Your Sponsored Products campaigns generate a search terms report showing the actual queries shoppers typed before clicking your ad. This data bridges your marketing spend to actual customer behavior.

The actual search term. Not the broad-match keyword you targeted, but the specific words the customer typed. This is invaluable for understanding customer language and intent — and for discovering which intent clusters are generating your paid traffic.

Attributed sales. Revenue directly generated from clicks on this search term through your ad. This is the "what did I get?" number.

Cost per click (CPC). What you paid for each click. This varies by keyword competitiveness and time of day. It is the "what did I pay?" number.

Conversion rate. The percentage of ad clicks that turned into purchases. This combines ad relevance, listing quality, and price competitiveness into a single number.

Impressions and clicks. How often your ad was shown and how often it was clicked. The ratio is your ad click-through rate — a signal of ad relevance and creative quality.

When you map your PPC search terms back to your intent clusters, patterns emerge immediately. Some intents are generating excellent paid ROI. Others are burning budget on low-conversion terms. The overlap between PPC data and intent-level SQP performance is where the most actionable decisions are made.

2.4 THE MYTH OF THE GOLDEN RATIO

Many industry experts still preach the "Golden Ratio" of Amazon traffic: roughly 66% organic to 33% PPC. The prevailing theory is that PPC should be used primarily to stimulate organic ranking, which then subsidizes the cost of paid traffic.

However, the data tells a different story. In the highly competitive modern Amazon landscape, many highly successful and profitable businesses operate on the exact opposite ratio: 33% organic and 66% PPC.

Why? Because Amazon has fundamentally changed its search results page. The top of the page is now heavily saturated with Sponsored Products, Sponsored Brands, and video ads. PPC is no longer just a catalyst for organic rank — it is the permanent rent you pay for top-shelf visibility. If

your unit economics (TACoS and customer lifetime value) are healthy, artificially restricting your PPC volume just to chase a 66% organic ratio will choke your absolute profit.

The intent-centric Growth Machine optimizes for total profitable volume at the intent level — not for an arbitrary traffic ratio. The question is never "what percentage of my traffic is organic?" The question is "what is the most profitable way to capture this intent, and can my margins support the investment?"

2.5 INTENT CLUSTERS — THE ORGANIZATIONAL LAYER

Intent clusters are the fourth data layer, and they transform how you read the other three.

Instead of asking "how am I performing on 'orthopedic neck pillow'?", you ask "how am I performing across all 25 keywords that express the 'Pillow for Neck Pain' intent?" NeonPanel's intent analytics aggregate SQP data across all the terms in each cluster, giving you a single intent-level view of your competitive position.

Without intent clusters, you see fragmented data: orthopedic pillow at 5% purchase share, memory foam pillow at 9%, cervical support pillow at 3%. Three separate data points suggesting modest, scattered performance.

With intent clusters, you see a unified strategic picture: "Pillow for Neck Pain intent, 247,000 monthly searches, your aggregate purchase share 6%, Competitor A at 22%." One number that reveals the size of the gap and the scale of the opportunity.

Ask the AI: "Show me my overall performance across all search terms in the Pillow for Neck Pain intent for the last 30 days. I want to see total impressions, impression share, click share, purchase share, and how I'm classified — visible but converting poorly, invisible entirely, or competitive."

The result replaces 25 rows of a spreadsheet with a single, actionable diagnosis.

2.6 THE RYG FRAMEWORK IN INTENT CONTEXT

To make Brand Analytics actionable, you need a classification system. The RYG (Red-Yellow-Green) framework provides exactly that — and in this edition, it operates at the intent level rather than the keyword level.

🟢 **Green Intent — Purchase Share ≥ 12%.** You own a meaningful portion of this intent's market. Shoppers regularly buy from you when searching this intent. Strategy: Defence — maintain share, monitor competitors, protect against erosion.

🟡 **Yellow Intent — Purchase Share 5–11%.** The opportunity zone. The market is large, you are visible, but you are not converting as effectively as you could. Strategy: Expansion — improve persuasion, increase visibility, gain share.

🔴 **Red Intent — Purchase Share below 5%.** The market is buying on this intent, but you are effectively invisible or losing every comparison. Strategy: Investigation — diagnose the root cause (indexing gap, listing quality, price), then target accordingly.

The thresholds are not sacred. For a highly fragmented category, Green might be set at 8%. For a category with a dominant player, 20%. Ask the AI to retrieve your current thresholds and update them for your market's competitive structure.

THE FOUR-LAYER SYNERGY

The power is not in any single layer — it is in the overlay.

SQP tells you where the market is buying. It reveals the keywords generating the most purchases and your share of each.

SCP tells you where your funnel leaks. It diagnoses whether you have a visibility problem, a click-through problem, a consideration problem, or a conversion problem.

PPC Search Terms tell you what you are paying for and whether it is working. It reveals the financial reality behind your advertising — which keywords generate ROI and which are burning cash.

Intent Clusters tell you which battlefield to fight on. They group all four layers into coherent strategic territories, revealing where you should concentrate, where you should defend, and where you should walk away.

Cross-referencing all four produces the intelligence that drives every decision in this book:

- An intent with high market size, low brand share, and zero PPC spend is a strategic gap — an untapped opportunity worth coordinated investment.
- An intent where your organic share is strong and your PPC spend is high is a cannibalization risk — you may be paying for clicks you would earn anyway.
- An intent where SCP shows a click-to-cart leak across all your ASINs is a listing quality problem — spending more on this intent will not fix it.
- An intent with low search volume, high CPC, and declining momentum is a drain — a candidate for budget reallocation to higher-opportunity intents.

This four-layer framework is the analytical engine of every chapter that follows.

Chapter 3: Setting Up Your Intent Machine

CONNECTING YOUR DATA, DEFINING YOUR STRATEGY

The concepts in this book are universal. Any brand-registered Amazon seller can manually pull SQP data, download PPC reports, and build spreadsheet analyses. Many do. The problem is that manual analysis takes hours, creates stale data, and does not scale.

The intent-centric Growth Machine uses an AI agent connected to a structured data layer. This is not about replacing your strategic judgment. It is about replacing the spreadsheet.

WHAT MCP IS (FOR NON-TECHNICAL READERS)

The Model Context Protocol is an open standard that allows an AI assistant to call structured tools through natural-language conversation. Instead of you logging into Seller Central, navigating to Brand Analytics, downloading a CSV, opening it in Excel, and building a pivot table, you simply describe what you want:

```
"Show me the top 20 keywords for my donut pillow by purchase share, and flag any
where my share is below 10%."
```

The AI translates this into a tool call, executes it against your synchronized data, and returns the result in seconds — already filtered, classified, and ready for analysis. You focus on the question and the decision. The machine handles the data retrieval.

THE NEONPANEL DATA LAYER

NeonPanel is the specific platform we use throughout this book. It connects to your Amazon Seller Central and Advertising accounts, synchronizes reports on a regular cadence, and exposes them through AI-friendly tools.

The data it unifies:

- **Brand Analytics:** Search Query Performance, Search Catalog Performance, Top Search Terms, Repeat Purchase behavior.
- **Advertising:** Campaign performance, search term reports, budget and bid data.
- **Financials:** Revenue, fees, cost of goods (landed cost and FIFO), profit and loss.
- **Inventory:** Stock levels, warehouse balances, sales velocity, replenishment risk.
- **Supply Chain:** Lead times, PO recommendations, shipment tracking, arrival forecasts.

• **Intent Layer:** User-defined intent clusters, AI-assisted clustering, intent performance analytics.

This is the nervous system of the Growth Machine. Every analysis in this book draws from this unified layer.

SETUP: SIX STEPS TO YOUR FIRST INTENT ANALYSIS

Step 1: Connect Your Accounts

Integrate your Amazon Seller Central and Advertising accounts through NeonPanel's integration flow. Once connected, verify everything is working by asking the AI:

```
"List my connected companies and confirm which one is active."
```

This returns your company identifier, which is used in every subsequent analysis.

Step 2: Map Your Product Catalogue

Ask the AI:

```
"Show me the full ASIN catalog for my account, including each parent ASIN and
all its child variations."
```

This is the foundation: a parent ASIN with five color variations needs all five children included when you pull SQP data, or your purchase share will be understated.

Step 3: Configure Your RYG Thresholds

Review your market's competitive structure, then:

```
"Show me my current Red-Yellow-Green thresholds, and update them so that Green
starts at 8% purchase share — my category is fragmented and no seller holds more
than 12%."
```

The AI retrieves your current thresholds and saves your calibration.

Step 4: Bootstrap Your Intent Clusters

This is the step that separates intent-centric analysis from keyword-level analysis.

Option A — AI-Assisted Clustering (recommended for first setup):

> "Here are my top 100 search terms from the last 30 days. Please cluster them into 8 to 12 intent groups based on underlying customer goals. For each proposed intent, give it a descriptive name, explain what the customer is trying to achieve, list the search terms that belong to it, and score your confidence in the grouping."

The AI proposes 8–12 intent clusters. Review them carefully — this is a 10-minute investment that shapes every analysis going forward. Refine any clusters that feel too broad, too narrow, or misaligned with how your customers actually think. Then ask the AI to save the clusters you accept.

Option B — Manual Clustering (more control):

Think about your product category. What are the distinct customer goals? For a pillow seller, they might be: Pillow for Pain Relief, Pillow for Side Sleepers, Cooling Pillow for Hot Sleepers, Pillow for Allergies, Premium or Designer Pillow, Budget Pillow. For each one, describe the intent and list the search terms that belong to it. The AI saves each cluster as you define it.

Step 5: Verify and Refine

Ask:

> "List all my intent clusters. For each one, show me how many search terms it contains and flag any terms that appear ambiguous — ones that could reasonably fit in more than one intent."

Review each cluster for clarity. Does the name describe a real customer goal? Are the mapped terms genuinely related? Is the term count reasonable — typically 5 to 30 terms per intent? If one cluster contains 100 terms, it is too broad and should be split.

Intents are not permanent — you will refine them as you learn more about your market.

Step 6: Run Your First Intent Analysis

Ask:

> "Analyze my Pillow for Neck Pain intent over the last 30 days. I want total searches across all terms in the cluster, my brand's impression share, click share, and purchase share aggregated across the whole intent, an opportunity classification, and which individual search terms within the intent are driving the most conversion for me."

This moment is like the first keyword analysis in the original Growth Machine — except instead of one row in a spreadsheet, you see the entire strategic territory at once. Intents that looked like modest performance at the keyword level often reveal either hidden scale or hidden threats when aggregated.

A NOTE ON TOOLING AGNOSTICISM

While this book uses NeonPanel extensively, the analytical frameworks are tool-agnostic. You can apply the Intent Gap Analysis, the Incrementality Framework, and the Marketing Throttle with any data source — including manual CSV downloads from Brand Analytics.

The tools make it faster. The thinking is what matters.

That said: speed matters more than most sellers realize. A seller who runs this analysis in 15 minutes every Monday will outperform a seller who runs it once a quarter. In a competitive market, 52 learning cycles per year beats 4 learning cycles per year. Tooling is a competitive advantage when it compresses the loop.

SECTION II: INTENT-DRIVEN MARKETING & ADVERTISING

Chapter 4: The Intent Gap Analysis

FINDING HIGH-VOLUME INTENTS WHERE YOUR BRAND SHARE IS BELOW 10%

If you take only one concept from this entire book, take this one.

The Intent Gap is the distance between where the market is buying — at the intent level — and where your brand is capturing share. It is the single most actionable output of the intent-centric Growth Machine, because it answers the question every seller should be asking but few articulate clearly: where are customers searching, finding, and buying products like mine — but not buying mine?

Every intent where customers are purchasing and your brand is absent or underrepresented is a gap. Every gap is revenue you are leaving on the table. Not theoretical revenue — actual purchases happening right now, going to your competitors, on intents you could compete for.

WHY INTENT GAPS EXIST

Intent gaps exist for the same reasons keyword gaps exist — but at a higher, more strategic level. The reason determines the response.

You are not indexed on any variant of this intent. Your listing does not contain language that connects to the customer goal. A pillow seller whose title reads "Memory Foam Pillow — Blue" may not appear on any "neck pain" search because the listing never mentions the benefit. Amazon's algorithm cannot connect your product to the intent.

Response: Add intent-aligned language to your listing. Update your title, bullet points, and A+ content to speak directly to the customer goal. This is the simplest and highest-ROI fix.

You are indexed but underperforming across the intent. You appear on 10 of the 27 keywords that express "Pillow for Neck Pain," but your average purchase share across those terms is only 3%. You have a presence, but it is not translating to market share.

Response: Diagnose the funnel leak. Are you being shown (visibility problem)? Clicked (listing quality problem)? Purchased (conversion problem)? The answer determines the fix.

You are indexed on individual keywords but have not invested to grow them together. You have organic presence on "orthopedic neck pillow" at 5% purchase share and "cervical support pillow" at 4%, but you are running isolated campaigns, not a coordinated intent strategy.

Response: Launch a unified PPC campaign targeting the entire intent, not isolated keywords. Consolidate your bidding strategy around the customer goal rather than individual terms.

You are underestimating the intent's size. You are bidding on "orthopedic pillow" (15K searches per month) without realizing that the broader "Pillow for Pain Relief" intent encompasses 250K searches per month across 27 keyword variants. You are fishing in one creek while the river runs nearby.

Response: Expand your target to the full intent. Pull all related terms. Allocate budget to match the true opportunity size.

THE INTENT GAP ANALYSIS: FIVE STEPS

Step 1: Pull Aggregated Intent Performance Data

Ask the AI:

> `"Show me my performance across all my defined intents for the last 30 days. For each intent, give me the total search volume across all terms in the cluster, my brand's purchase share, and the Red-Yellow-Green classification."`

This is your complete map of intent-level competitive position — the starting point for every strategic decision.

Step 2: Filter for the Gap

From the full intent list, isolate the ones matching two criteria:

High Market Size. Intents with 50,000 or more monthly searches. The exact threshold depends on your category — in a niche market, 10,000 searches might be significant. Use your top quartile of intents by volume as the starting filter.

Low Brand Share. Intents where your purchase share is below 10% — or whatever you have calibrated as your Red or Yellow threshold. These are intents where the market is buying but you are not capturing a meaningful portion.

The intersection of these two filters is your Intent Gap List: high-opportunity, low-performance intents. These are your strategic targets.

Step 3: Diagnose the Funnel

For each gap intent, ask the AI:

> "Break down my funnel performance on the Pillow for Neck Pain intent, by each of my ASINs. I want to see impression share, click share, and conversion share for the intent as a whole and for each individual product."

This tells you precisely where you are losing shoppers:

- Low impression share means an indexing or organic rank problem. You are not being shown.
- Low click share despite adequate impressions means an above-the-fold listing issue. You are being shown but not chosen.
- Low conversion share despite healthy clicks means a product detail page, price, or reviews issue. Shoppers are visiting and leaving.

Each diagnosis points to a different intervention. There is no point launching PPC on an intent where your listing's click-through rate is below average — you will pay for traffic that does not convert. Fix the listing first.

Step 4: Cross-Reference Momentum

Not all gap intents are worth pursuing equally. Ask the AI:

> "For the Pillow for Neck Pain intent, show me the search volume trend over the last 8 weeks. Are overall searches for this intent growing, stable, or declining? And which individual terms within the intent are growing fastest?"

Accelerating intents are priority targets — the market is moving toward them, which means early investment compounds. Stable intents are reliable opportunities. Declining intents should be deprioritized unless you see a clear seasonal pattern.

Also look for emerging terms within each intent — keywords with low current volume but sharp growth curves. These are the next wave. Identifying them early means cheaper CPCs and easier ranking before the rest of the market catches on.

Step 5: Check Competitive Density

Before committing budget to a gap intent, understand who currently owns it. Ask the AI:

> "For the Pillow for Neck Pain intent, show me the top three competitors by purchase share and tell me how dominant the market leader is — is this intent fragmented, does it have a weak leader, or is it dominated by one player?"

This produces a contestability picture for each intent:

- **Fragmented:** No single product holds more than 15% purchase share. The market is dispersed. You can break in with moderate investment.
- **Weak leader:** One product leads with 15 to 30%, but the remaining share is spread among many. Challengeable with a differentiated approach.
- **Dominant leader:** One product holds 40% or more. Breaking through is expensive and uncertain. Only pursue if this intent is strategically essential or if you have a genuinely differentiated product.

Contestability does not mean you avoid competitive intents entirely. It means you prioritize your limited budget on intents where the chance of meaningful share gain is highest.

OUTPUT: YOUR INTENT GAP REPORT

The synthesis is a prioritized intent list:

Intent	Market Size	Your Share	Gap Target	Momentum	Competitive Density	Recommended Action	Priority
Pillow for Neck Pain	247K	7%	12–15%	Growing	Weak leader	Launch unified PPC + listing update	1
Pillow for Side Sleepers	185K	11%	15%+	Stable	Fragmented	Increase bids, defend position	2
Cooling Pillow	92K	3%	8–10%	Accelerating	No clear leader	Listing update + targeted PPC	3
Budget Pillow	156K	9%	—	Declining	Leader at 45%	Low priority, monitor	—

This report becomes your strategic targeting roadmap. It feeds directly into PPC campaign strategy, listing optimization, and A+ content creation. It is not static — you regenerate it weekly as part of the Governance Loop and track how your share moves on each target intent over time.

Chapter 5: Intent-Level Advertising Incrementality

PROVING WHETHER YOUR PPC IS CANNIBALISING OR GROWING MARKET SHARE

Every Amazon seller running advertising faces the same nagging question: am I paying for clicks I would have gotten for free?

The question is not paranoia. It is a legitimate strategic concern backed by a simple economic reality. If your product already ranks in the top three organically on a high-volume intent, and you are also running Sponsored Products heavily across all 25 keywords in that intent, a meaningful percentage of the shoppers clicking your ads would have clicked your organic listing anyway. You are paying for nothing.

The flip side is equally important. If your product is invisible organically on a high-volume intent, and PPC is the only reason shoppers see your listing, then every ad click is genuinely incremental. Continuing that spend would directly mean driving growth. Cutting that spend would directly cost you sales.

The difference between these two scenarios is the difference between wasting money and investing it. The Growth Machine gives you the framework to tell them apart — now at the intent level rather than keyword by keyword.

THE THREE SCENARIOS

Scenario 1: The Waste Intent

An intent where you dominate organically but are also spending heavily on ads.

The telltale pattern: SQP shows high brand purchase share on the intent — 40% or more. Your organic rank on the intent's primary keywords is in the top three positions. Yet you are running Sponsored Products campaigns across 15 or more keywords within this intent.

What is happening: shoppers would find and choose your product anyway. The ads are a redundant signpost on a road that already leads to your door.

The nuance: this does not always mean you should pause entirely. There is a defensive argument — if a competitor is actively bidding on the intent, you may need a minimal advertising presence

to protect your position. The question is whether the cost of defence is proportional to the risk. If you are spending $1,800 per month defending an intent where you hold 68% organic purchase share, the economics likely do not justify it unless a competitor is specifically attacking.

Action: Reduce bid intensity on this intent. Redirect the liberated budget to gap intents identified in Chapter 4. Maintain a light defensive position if a competitor is actively present.

Scenario 2: The Growth Intent

An intent where your organic presence is weak but PPC is driving meaningful sales.

The telltale pattern: SQP shows low brand purchase share on the intent — below 8%. PPC-attributed sales on the intent's keywords are among your top performers. Organic rank for the intent's primary terms is on page two or lower.

What is happening: PPC is your only visibility on this intent. Every ad click is genuinely incremental — these are sales you would not have made without advertising.

Action: Maintain or increase spend. This intent is a growth investment, not a cost. Monitor closely whether sustained sales velocity from PPC begins building organic rank — the flywheel effect described next.

Scenario 3: The Flywheel Intent

The ideal outcome. An intent where you invested in PPC, gained sales velocity, and that velocity began pulling your organic rank upward.

The telltale pattern: historically, SQP showed low brand share on this intent. You launched aggressive PPC. Over four to eight weeks, brand purchase share in SQP began rising. Your organic impressions and clicks on the intent's keywords increased. Organic rank improved.

What is happening: the flywheel is spinning. PPC-driven sales velocity signalled to Amazon's algorithm that your product is relevant for this intent, and the algorithm rewarded you with improved organic rank. Now organic traffic is compounding on top of your paid traffic.

Action: This is the point where you begin gradually reducing PPC spend — not cutting it off, but easing it as organic takes over. The key metric to watch is the crossover point: the week when your organic purchase share on this intent exceeds your paid contribution. After that, PPC becomes maintenance and you redirect budget to the next growth intent.

THE INCREMENTALITY SCORE

To move beyond qualitative scenarios to a quantitative metric, calculate an Incrementality Score for each intent:

Intent Incrementality = (Paid Purchase Share ÷ Total Brand Purchase Share) × (1 − Organic Strength Factor)

Where Paid Purchase Share is purchases attributed to ads across all keywords in the intent, Total Brand Purchase Share is your overall brand share on the intent from SQP, and Organic Strength Factor is your organic purchase share as a fraction of total.

Interpreting the score:

- **Score near 1.0:** PPC is almost entirely incremental. Without ads, you would have no presence on this intent. This is a Growth intent.
- **Score near 0.5:** PPC is contributing, but you also have organic traction. This is a Flywheel intent — monitor the trend toward organic takeover.
- **Score near 0.0:** PPC is largely redundant. Your organic presence captures most of the traffic anyway. This is a Waste intent.

Ask the AI to calculate this for each of your intents by combining your SQP purchase share data with your PPC attributed sales data for the same keywords.

THE INTENT INCREMENTALITY MATRIX

For your weekly analysis, maintain a matrix of all intents:

Intent	SQP Purchase Share	PPC Spend Monthly	Incrementality Score	Scenario	Action
Pillow for Neck Pain	7%	$3,200	0.89	Growth	Maintain spend
Pillow for Side Sleepers	11%	$2,100	0.52	Flywheel	Monitor organic crossover
Brand + Pillow	68%	$1,800	0.12	Waste	Reduce bids significantly
Cooling Pillow	3%	$0	n/a	Gap	Launch campaign

This matrix connects directly to your Gap Report from Chapter 4 (intents with zero PPC spend and organic traction are the fastest wins) and informs every budget reallocation decision in the weekly governance loop.

Chapter 6: Intent Discovery and Blue Oceans

FINDING HIGH-VALUE INTENTS THAT YOUR COMPETITORS ARE MISSING

The most profitable intents in your category might be the ones nobody is aggressively competing for yet — not because they lack volume or purchase intent, but because competitive attention is concentrated on a few obvious intents.

Most sellers build their campaigns by mining their own PPC search terms report: they look at what is already working, add close variants, and incrementally expand. This is like fishing in a pond you already know. You will catch some fish, but you will never discover the lake over the hill.

Intent discovery finds the lake.

THREE DISCOVERY METHODS

Method 1: Visible But Untargeted Intents

These are intents where you have some organic presence — SQP shows non-zero purchase share — but you are running zero PPC campaigns on the intent's keywords.

Ask the AI:

> `"Which of my defined intents show some organic purchase share in SQP – even if it's small, say above 2% – but have zero ad spend across all keywords in the cluster?"`

These intents appear because your listing naturally contains intent-aligned language. Amazon's algorithm associates your product with the customer goal, and some organic traffic flows. But you have never amplified it with advertising.

This is the fastest win available. You are already indexed. You just need to add the intent to your campaigns and let your existing relevance do the work.

Method 2: Category Universe Gaps

These are intents that are active in your category's total market — appearing in Amazon's Top Search Terms or in competitor SQP data — but are entirely absent from your own SQP results.

The market cares about the intent. You are not even in the conversation.

Ask the AI:

> "Look at the top search terms for my category. Which ones don't appear in my SQP data at all — meaning Amazon currently has no association between my products and those searches?"

These require foundational work first: update your listing content to address the intent, verify indexing, then test PPC once you have confirmed visibility.

Method 3: Efficiency Signals

These are intents where your purchase share is already healthy, but your cost per click is unusually low compared to your portfolio average. The market is buying on this intent, competition is not bidding aggressively, and you are getting cheap access.

Ask the AI:

> "Which of my active intents have a healthy purchase share — above 6% — but a cost per click that's significantly below my portfolio average? I'm looking for intents where I'm winning relatively cheaply."

This is an arbitrage opportunity. Increase budget allocation to these intents before competitors notice. Monitor CPC trends closely — if others start bidding, the arbitrage window closes.

BUILDING THE INTENT DISCOVERY REPORT

The synthesis of all three discovery methods produces your Blue Ocean Map:

Discovery Source	Intent	Market Size	Your Current Share	Priority	Rationale
Visible, Untargeted	Pillow for Hot Sleepers	68K	2.1%	High	Already indexed organically; launch PPC to amplify
Efficiency Signal	Pillow for Stomach Sleepers	41K	9%	High	Healthy share, low CPC — scale before competitors notice
Category Gap	Pillow for Pregnancy	92K	0%	Medium	Not indexed yet; listing update first, then test PPC
Category Gap	Hypoallergenic Pillow	55K	0%	Medium	Market cares; you are entirely absent

This report feeds directly into campaign expansion decisions, listing optimization priorities, and budget allocation. Updated weekly, it is the primary source for new campaign launches.

Chapter 7: True Unit Economics for Intent Strategy

USING LANDED COST TO SET YOUR MAXIMUM AGGRESSIVE BID PER INTENT

Everything in Chapters 4 through 6 assumed your business can afford to compete. This chapter confronts the constraint that determines how aggressively you actually can: your true unit economics.

Amazon Seller Central shows you revenue and fees. Your advertising console shows you ad spend. Neither knows your manufacturing cost, your sea freight rate, your customs duties, your inspection fees, or your inland logistics. Until you unify all three layers — revenue, Amazon costs, and true cost of goods — your profitability number is fiction. And making growth decisions based on fictional margins is how sellers scale themselves into financial trouble.

THE PROBLEM WITH AMAZON'S PROFITABILITY VIEW

Seller Central's reports show you ordered units, revenue, and a rough breakdown of FBA fees and commissions. If you subtract ad spend from revenue minus fees, you get something that looks like profit.

It is not profit.

It is revenue minus the costs Amazon knows about. It excludes:

- **Manufacturing cost per unit.** What you actually paid your factory.
- **Freight.** Sea freight, air freight, or courier, allocated per unit.
- **Customs duties.** Import tariffs specific to your product's HS code.
- **Inspection fees.** Pre-shipment quality control, allocated per batch.
- **Insurance.** Marine cargo insurance, typically 0.3 to 0.5% of shipment value.
- **Inland transport.** Factory to port, port to warehouse, warehouse to FBA.
- **Labelling and prep.** If you use a prep centre, the per-unit handling fee.

For a typical private-label product with a $25 selling price, these excluded costs might total $6 to $8 per unit. Ignoring them overstates your margin by 24 to 32%. That error compounds across every intent strategy decision you make.

THE TRUE CONTRIBUTION MARGIN FORMULA

True Contribution Margin = Selling Price − FBA Fees − Referral Commission − Landed Cost

Where Landed Cost encompasses every cost from factory floor to Amazon's fulfilment centre, allocated per unit:

Landed Cost = Manufacturing + Freight + Duties + Insurance + Inspection + Inland Transport + Prep

The True Contribution Margin is what remains before advertising. It is the maximum you can spend on acquiring a customer on this product and still break even. Everything spent below this margin is profit. Everything above it is a deliberate investment in market share — which may be strategic, but must be conscious.

Ask the AI:

> "Pull my landed cost per SKU and combine it with my revenue and FBA fee data to show me the True Contribution Margin for each product. I want to see exactly how much remains per unit before I account for advertising spend."

SETTING YOUR MAXIMUM AGGRESSIVE BID

With True Contribution Margin in hand, the advertising question transforms from "what CPC feels reasonable?" to "what CPC does my unit economics support?"

The formula:

Maximum Break-Even CPC = True Contribution Margin × Conversion Rate

Where Conversion Rate is the percentage of ad clicks that result in a purchase, pulled from your PPC data.

Example for a mid-market pillow:

- Selling Price: $29.99
- FBA Fees: $6.50
- Referral Commission: $4.50
- Landed Cost: $7.20
- **True Contribution Margin: $11.79**
- Conversion Rate: 14%
- **Max Break-Even CPC: $11.79 × 0.14 = $1.65**

You can bid up to $1.65 per click on any keyword within this product's intents and break even on every resulting sale. Bidding $1.00 gives you $0.65 profit per ad-driven sale. Bidding $2.00 means you are investing $0.35 per sale into market share growth — acceptable as a deliberate short-term strategy, unsustainable as a permanent position.

This number is your financial guardrail. Every intent campaign recommendation in the Growth Machine is constrained by it.

INTENT-LEVEL BID STRATEGY ACROSS MULTIPLE ASINS

An intent may have multiple ASINs competing within it. Your "Pillow for Neck Pain" intent might include a premium pillow at $45 with an $11 margin, a mid-market pillow at $28 with a $6.50 margin, and a budget option at $18 with a $3.50 margin.

Different ASINs have different maximum break-even CPCs. Ask the AI:

> "For each of my active ASINs, calculate the True Contribution Margin and the maximum break-even CPC given the current conversion rate. I want to see which products can afford to compete aggressively for the Pillow for Neck Pain intent and which ones cannot."

The result clarifies your strategy:

- **Premium ASIN:** Can bid up to $1.59, so set bids at $1.20 to $1.40 to capture share while maintaining a margin buffer.
- **Mid-market ASIN:** Can bid up to $0.67, so set bids at $0.50 to $0.60.
- **Budget ASIN:** Cannot afford to bid aggressively. Either run minimal bids or exclude from this intent's campaigns entirely.

This prevents the common mistake of bidding equally across all your products and watching your margins erode on your low-price items.

FIFO AND BATCH-SPECIFIC MARGINS

There is a subtlety most profitability models ignore: your cost of goods is not static. Each inventory shipment arrives with a different landed cost. Sea freight rates fluctuate. Factory prices change. Exchange rates move.

FIFO (First In, First Out) cost analysis addresses this. Ask the AI:

> "Show me my FIFO cost analysis. I want to see which inventory batches are currently selling, what each batch cost per unit, and how my effective cost of goods will change as each batch sells through."

A high-cost shipment temporarily depresses your margin until it sells out and cheaper inventory takes over. This matters for intent strategy:

- **Bid timing:** Bid more conservatively on growth intents while high-cost inventory is selling, then ramp aggressively when cheaper stock arrives.
- **Pricing decisions:** If your current batch has unusually high COGS, a modest price increase may be more profitable than absorbing the margin hit.
- **Profit tracking:** Your weekly P&L should reflect the actual FIFO cost for the units sold, not a blended average.

Chapter 8: The Marketing Throttle

ENSURING INTENT GROWTH DOES NOT COLLIDE WITH INVENTORY REALITY

The intent-centric Growth Machine is powerful at identifying opportunities. It can show you that the "Pillow for Neck Pain" intent is growing, your market share is low, and you could grow revenue by 40% by scaling bids.

Then you run out of inventory.

A stockout is not just a missed sale. It is a compound loss:

- **Immediate revenue loss.** Every day out of stock is a day of zero sales on that ASIN.
- **Organic rank erosion.** Amazon's algorithm weighs recent sales velocity heavily. A stockout craters your velocity, and your organic rank drops. When you restock, you are climbing back from a lower position than when you left.
- **Competitor absorption.** Your traffic does not disappear — it goes to the next product in search results. Your competitor's velocity increases, their rank improves, and they occupy the position you vacated. You are now competing against a stronger opponent when you return.
- **Review velocity stall.** No sales means no new reviews. In competitive categories, a two-week review drought is noticeable.
- **Flywheel destruction.** The worst outcome: you invested heavily in PPC to build momentum on a key intent, your velocity spiked, your organic rank rose — and then you stocked out. The

PPC spend is gone. The organic rank decays. The competitor who was below you has overtaken you. You have to restart the flywheel from scratch.

The Marketing Throttle prevents this. It is a hard rule: no bid increase, no budget expansion, no new intent campaign launch is ever approved without first checking inventory health.

THE THROTTLE TABLE

Days of Supply	Mode	What It Means	Intent Actions Permitted
More than 45 days	🟢 **Aggressive**	Growth window. Full budgets, competitive bids, new campaign launches.	Launch campaigns on new gap intents. Increase bids on growth intents. Test untargeted intents.
20 to 45 days	🟡 **Maintenance**	Adequate but tightening. Hold current campaigns steady.	Maintain current intent campaigns. No new intent launches. No bid increases. Focus on efficiency.
Less than 20 days	🔴 **Conservation**	Danger zone. Reduce bids 30%. Pause all new initiatives.	Pause all discovery and broad-match campaigns. Keep only highest-converting intent keywords running.
Less than 10 days	⛔ **Emergency**	Critical. Preserve every unit for organic sales.	Pause all advertising except branded terms. Accept lower growth. Wait for replenishment.

The throttle operates per ASIN. When you run intent campaigns targeting multiple ASINs, each product gets its own throttle mode. An intent campaign might continue at full strength for ASIN1 (45 days of supply) while being paused entirely for ASIN2 (12 days of supply), even though both ASINs appear in the same intent cluster.

CHECKING INVENTORY HEALTH

Ask the AI:

> `"Show me current inventory health for all my active SKUs. For each one, I need the days of supply based on current sales velocity, the throttle mode it should be in, and whether any replenishment orders are already in transit."`

Three data points combine to give you the full picture:

- **Days of supply** based on current velocity — not the velocity from two months ago when you last updated the model.
- **Inbound shipments** already in transit, which will relieve pressure within days or weeks.
- **Purchase order status** for any replenishment not yet shipped, which tells you whether the stockout is avoidable or inevitable.

THE LEAD-TIME BUFFER

The Marketing Throttle tells you when to slow down. The Lead-Time Buffer tells you when you needed to act to avoid slowing down at all.

Every product has a total lead time — the number of days from placing a purchase order to receiving inventory in Amazon's warehouse, available for sale:

Total Lead Time = **Manufacturing Time** + **Transit Time** + **Receiving Time** + **Safety Buffer**

Manufacturing time is typically 15 to 30 days for standard products. Sea freight transit adds 25 to 40 days. Amazon's receiving process currently adds 7 to 14 days. Add a safety buffer of at least 7 days for unforeseen delays. For most private-label sellers, total lead time is 55 to 90 days.

Last Possible Launch Date = **Estimated Stockout Date** − **Total Lead Time**

If your stockout date is June 1 and your total lead time is 60 days, you needed to place a purchase order by April 2. If today is April 10 and no order has been placed, you will stock out, and the only option is to enter Conservation Mode and start expediting.

Ask the AI:

> `"Based on current velocity and inventory levels, what is the estimated stockout date for each of my SKUs, and what is the last possible date I can place a purchase order and still avoid a stockout assuming my standard lead times?"`

This analysis, run weekly, turns stockouts from surprises into choices.

THE THROTTLE IN THE GOVERNANCE LOOP

The Marketing Throttle check is the first step in every weekly governance ritual — not the last. Before you look at intent opportunities, before you review keyword gaps, before you evaluate incrementality scores, you check inventory health for every SKU. The throttle mode each product is in determines the boundary conditions for all decisions that week.

A beautifully constructed intent gap analysis is worthless if you are 14 days from stockout.

SECTION III: OPERATIONS & PROFITABILITY

Chapter 9: Stockout Prevention and Intent-Aware Inventory Management

The Marketing Throttle from Chapter 8 governs individual ASINs. Intent strategies span multiple ASINs. That intersection requires nuance.

When you run the throttle check during your weekly governance loop, you assign each ASIN its mode independently. But a "Pillow for Neck Pain" intent campaign might target three or four of your ASINs simultaneously. If ASIN2 is in Conservation mode while ASIN1 and ASIN3 are in Aggressive mode, your intent strategy must reflect those differences:

- Continue or expand campaigns targeting ASIN1 and ASIN3, which have healthy inventory.
- Hold ASIN2's existing ads but make no bid increases or budget expansions.
- If ASIN2 hits Emergency mode, pause its ads entirely and shift all intent campaign focus to ASIN1 and ASIN3.

Ask the AI:

> `"For each of my intent clusters, show me the inventory throttle mode for every ASIN that appears in those campaigns. Tell me which ASINs are cleared for growth actions and which are constrained."`

This gives you a clean per-intent, per-ASIN action matrix each Monday before you make any campaign decisions.

The deeper principle: intent strategies amplify velocity, and velocity depletes inventory faster than baseline sales projections predict. When you successfully capture share on a high-volume intent, your days-of-supply calculation from last week may already be dangerously optimistic. Recalculate it every Monday using the most recent velocity data — not the number from when you last updated your planning model.

SECTION IV: THE WEEKLY GOVERNANCE LOOP

Chapter 10: The Weekly Intent Governance Ritual

TRANSFORMING DATA INTO A PERPETUAL MOTION MACHINE

You have the tools. You have the frameworks. You have the data pipeline. Now comes the part that separates the sellers who grow from the sellers who analyze once and then forget about it for three months.

The Growth Machine is not a one-time audit. It is not a quarterly review. It is a weekly ritual — a structured cadence that turns data into decisions, decisions into actions, and actions into measurable outcomes. The loop is the machine. Everything else is just components.

WHY WEEKLY?

Not daily. Not monthly. Weekly.

Why not daily? Amazon's data has a 48 to 72 hour reporting lag. Brand Analytics data is aggregated weekly. PPC data shows attributions over a 7-day or 14-day window. Reacting to daily fluctuations means reacting to noise, not signal. Daily adjustments create oscillation — bid up, bid down, bid up — that produces worse results than steady-state bidding.

Why not monthly? A month is too long to detect and respond to competitive shifts. If a new competitor enters your top intent and starts gaining purchase share, four weeks of unchecked erosion can cost you organic rank that takes months to rebuild. Weekly cadence is fast enough to catch problems and slow enough to see real trends.

Why weekly? Because the compounding effect is real. The seller who runs 52 learning cycles per year will outperform the seller who runs 12. Not by a little — by a lot. Each cycle teaches you something about your market that your competitor does not yet know. Over a year, that accumulates into a proprietary understanding of your category that no tool or shortcut can replicate.

THE FIVE-STEP RITUAL

Every Monday morning. Same time. Same sequence. Non-negotiable.

Step 1: Inventory Health Check — The Throttle

Before any marketing analysis, check inventory. Ask the AI:

> "Show me the current days of supply for every active SKU, based on the most recent sales velocity. Assign each one its throttle mode — Aggressive, Maintenance, Conservation, or Emergency."

Record the throttle mode for every SKU. This is your Week N baseline and the constraint layer for all decisions that follow.

For each intent, note which ASINs are in which throttle mode. An intent campaign targeting four ASINs might be in Aggressive mode for three of them and Conservation mode for one. Treat each ASIN independently.

If any SKU has moved from Aggressive to Maintenance since last week, investigate immediately. Did velocity spike? Did a restock get delayed? Is an inbound shipment in transit that will relieve the pressure?

Step 2: Review Last Week's Intent Actions

Open your action log from the previous week. For each action taken — each bid change, budget shift, new intent campaign, listing edit — ask: did the predicted outcome occur?

If you increased your bid on the Pillow for Neck Pain intent with the hypothesis that purchase share would increase by 2% within two weeks, check the SQP data. Did purchase share move? By how much? Did CPC stay where you expected, or did competition drive it higher?

For each action, record three things:

- **What happened:** The actual outcome versus your prediction.
- **The delta:** How far off was your hypothesis?
- **The learning:** What do you now know about your market that you did not know before?

This feedback loop is what transforms a static analysis into a learning system. Over time, you build institutional memory. "In our category, intent-level bid increases of $0.15 typically yield 1.5% purchase share gain within 10 days." That calibration is priceless — and it only comes from recording and reviewing outcomes week after week.

Step 3: Run the Full Intent Data Pull

Execute the Growth Machine data extraction for the current week. Ask the AI:

> "Pull this week's intent performance data for all my defined intents. Show me aggregated purchase share, click share, and impression share for each one. Then compare to last week's numbers and flag any intent where purchase share moved more than 2% in either direction."

Then follow up:

> "Show me a side-by-side comparison of my top five intents by purchase share this week versus last week. Which intents are improving and which are declining?"

And:

> "Run the intent opportunity ranking for this week and tell me if any new hidden gems, leaky intents, or emerging opportunities have appeared that were not in last week's report."

Compare to your Week N−1 baseline. Calculate deltas. Flag any significant change in purchase share, click share, or TACoS for investigation.

Step 4: Generate This Week's Intent Actions

With the delta analysis complete, generate a new action list. The most common triggers:

New gap intent identified. A search term cluster that was not in last week's gap report has appeared — perhaps a seasonal trend is emerging, or your AI spotted an intent you had not defined yet. Create the intent. Plan the listing update. Schedule the PPC launch.

Cannibalization detected. An intent where organic share has grown to the point that PPC is now mostly redundant. The incrementality score crossed below 0.15. Reduce bids. Redirect budget.

Flywheel accelerating. An intent where purchase share rose despite flat PPC spend. Organic rank is improving. Begin easing bids slightly as organic takes over. Monitor the crossover point.

Stockout risk escalation. A SKU moved from Aggressive to Maintenance. No new intent campaigns for that product this week. Check PO status. Calculate Last Launch Date.

Funnel leak found. SCP shows your click-through rate on a top intent dropped 15% this week. A competitor may have updated their main image or lowered their price. Investigate and plan a response.

For each action, record four things:

- **What:** The specific change ("Increase bid on Pillow for Hot Sleepers intent from $0.85 to $1.05 for ASIN1 only").
- **Why:** The data that prompted it ("Purchase share rose from 4% to 6% on this intent, incrementality score is 0.87, ASIN1 is in Aggressive mode").
- **Hypothesis:** The expected outcome ("Purchase share will reach 8% within 14 days").
- **Constraint check:** Inventory mode is Aggressive, True Contribution Margin supports a Max CPC of $1.59.

Step 5: Execute and Log

Implement the changes. Update bids. Adjust budgets. Launch campaigns. Update listings.

Log every action with a timestamp. Close the loop. Wait seven days.

THE COMPOUNDING EFFECT

The first iteration of the loop produces rough insights. The second iteration is slightly better because you have one week of hypothesis-versus-outcome data. By the eighth week, you have two months of calibrated learning about your specific market.

After a quarter, the Growth Machine is not just analyzing — it is predicting. You know that a certain intent responds to bid increases within 10 days. You know that another intent's purchase share is seasonal and peaks in September. You know that your main competitor typically drops prices when a new entrant enters their top keywords.

This accumulated intelligence is the true competitive advantage. The data is available to every brand-registered seller. The tools are available to anyone willing to set them up. But the compounded learning of 12 or 24 weekly cycles — that belongs only to you.

SECTION V: THE PERSISTENT BRAIN

Chapter 11: Intent Learning and Long-Term Memory

BUILDING THE MACHINE THAT NEVER FORGETS

A Growth Machine without memory is a Growth Session. You run the analysis, get insights, make decisions, and close your laptop. Next Monday, you start from scratch. The insights from last week are half-remembered. The hypothesis you tested has no recorded outcome. The competitive shift you noticed gets re-discovered three weeks later when it has already caused damage.

Memory transforms a session into an engine. An engine that accumulates experience, refines its models, and produces increasingly accurate predictions over time.

WHAT THE MACHINE NEEDS TO REMEMBER

The Action Log

Every action taken as part of the weekly governance ritual needs to be recorded with enough context to evaluate its outcome. Each entry should capture:

- **Timestamp.** When the action was taken.
- **Action type.** Bid change, budget change, intent campaign launch, keyword addition, keyword pause, listing update, pricing change.
- **Target.** The specific intent, keyword, campaign, ASIN, or SKU affected.
- **Before state.** The pre-change value — bid was $0.80, purchase share was 7%, PPC spend was $1,200 per month.
- **After state.** The post-change value — bid is now $1.05, intent campaign launched, listing updated.
- **Hypothesis.** The expected outcome and timeframe — "Purchase share on Pillow for Neck Pain intent will increase from 7% to 9% within 14 days."
- **Data context.** What the data showed that prompted the action — RYG classification, momentum score, incrementality score, competitive density, inventory throttle mode.

The Hypothesis Result Log

After the appropriate waiting period — typically 7 to 14 days depending on the action type — the system checks each open hypothesis:

- **Predicted outcome.** What you expected to happen.
- **Actual outcome.** What happened.
- **Delta.** The gap between prediction and reality — was your hypothesis 10% optimistic? 40% optimistic? Accurate?
- **Confidence adjustment.** Based on accuracy, calibrate how much weight to give similar recommendations in the future.
- **Causal notes.** Any confounding factors — a Prime Day sale skewed the data, a competitor ran out of stock, a new entrant appeared on the intent.

Over time, this log builds a calibration curve. If your bid-increase hypotheses are consistently 30% optimistic for your category, the machine learns to discount future predictions accordingly. If listing updates consistently deliver more than you projected, you learn to weight them more heavily in future planning.

The Weekly Snapshot Archive

Every Monday, the Growth Machine captures a snapshot of all key metrics:

- Purchase share per intent and per target keyword.
- Click share per intent — trending direction.
- TACoS — total advertising cost of sale, tightening or loosening.
- Days of supply per SKU — inventory throttle mode.
- Revenue and margin per product.
- Active intent campaigns and total ad spend.
- Any significant competitive events noticed during the data review.

The snapshot archive enables trend analysis and seasonality detection. When your purchase share on the Cooling Pillow intent dips in March, you can check whether it dipped in March last year too. When your CPC spikes in November across all intents, the archive shows the Q4 pattern and lets you plan budget accordingly — rather than reacting with surprise.

WHERE TO STORE IT

The memory layer has three tiers, from lightest to most persistent:

Tier 1: Session Memory

During a Growth Machine session, the AI agent maintains working context — the data it has pulled, the analysis it has run, the actions it has recommended. This disappears when the session ends. It provides within-session coherence but no cross-session learning.

Tier 2: NeonPanel's Historical Layer

NeonPanel already stores historical Brand Analytics data, financial statements, inventory records, and advertising reports. This is your read-only historical archive — the raw data that the Growth Machine analyzes each week. It persists automatically and provides the factual foundation for trend analysis.

Tier 3: Your Action and Learning Database

The Action Log and Hypothesis Result Log need a dedicated store that you control and update. Options include:

- **A structured spreadsheet** (Google Sheets or Excel) with columns matching the fields above. Low-tech, easy to maintain, searchable by intent, action type, or outcome. This is where most sellers should start.
- **A database** (Airtable, Notion database, or a simple tool like Supabase) for more structured querying. Useful once you have 50 or more logged actions and want to filter and analyze patterns across them.
- **A plain text file in your workspace** updated each week. The simplest option — create an action log document and append to it every Monday. Easy to share with a team or reference during strategy reviews.

The important thing is not the tool. It is the discipline. If you do not log the action and the hypothesis, you cannot evaluate the outcome. If you do not evaluate outcomes, the machine does not learn. If the machine does not learn, you are running the same analysis from scratch every week — slightly better informed by experience, but without the quantitative calibration that turns a good analyst into a precision instrument.

INTENT MATURATION

Over 12 to 24 weeks, your intent strategies mature in ways that cannot be replicated by any competitor who is still thinking at the keyword level.

You learn which intents respond predictably to bid pressure — typically fragmented intents with weak leaders where any increase in visibility directly translates to purchase share gain. You learn which intents are seasonal — growing in Q4, declining in Q2, with a predictable recovery curve. You learn which intents have inelastic demand — where price cuts do not improve share because brand loyalty is driving the decision, not price sensitivity. You learn which intents are genuine

entry opportunities — where competition is low, searches are growing, and your organic rank is quietly building week after week.

This accumulated knowledge is your competitive moat. Every other seller sees the same SQP data. But you understand the rhythm of your specific market. You have 24 weeks of calibrated hypotheses and measured outcomes. You have a model of your category that took months to build and cannot be copied.

That is the real output of the Persistent Brain.

SECTION VI: THE NEONASPHERA AGENT

Chapter 12: The Intent-Warm AI Agent

SYSTEM PROMPT AND DECISION CHAIN FOR INTENT-CENTRIC GROWTH

Everything in this book — the frameworks, the analysis methods, the governance ritual — can be performed manually. You can pull data, build spreadsheets, and make decisions entirely with your own brain and spreadsheet tools. Many successful sellers do exactly that.

The NeonaSphera Agent is for sellers who want to automate the analytical layer and reserve their cognitive bandwidth for strategic decisions. It is an AI orchestration layer that follows a strict protocol — pulling data, applying the frameworks from this book, checking constraints, and generating prioritized recommendations.

It is not a chatbot. It does not hallucinate strategy. It follows a defined decision chain, always grounded in data, always constrained by your operational reality.

It does not make changes to your campaigns automatically. It analyzes and recommends. You review and execute. This is deliberate. Automated execution sounds efficient but removes the human judgment that catches edge cases, understands seasonal context, and makes the strategic trade-offs that no algorithm can quantify.

The agent is your analyst. You are the CEO.

THE SYSTEM PROMPT

The system prompt defines the agent's identity, objective, and operating rules. Think of it as the job description for your AI analyst — the instructions that govern its behavior every single week, regardless of what data it encounters:

You are the NeonPanel Intent Growth Architect. Your objective is the systematic expansion of market share at the intent level while maintaining profitability for the brand you serve.

Operating Principles:

Data first. Every recommendation must cite the specific intent-level metric that supports it. Never speculate, never generalize. If the data does not clearly support a recommendation, say so and explain what additional data would be needed.

Constraints are non-negotiable. Never recommend spend increases or new campaign launches without first checking inventory health. If an ASIN is in Conservation or Emergency mode, no growth actions are permitted for that product — state this clearly and move on.

Intents are primary. Keywords are the underlying data layer, but intent-level aggregation is your primary analytical lens. Always interpret keyword-level data in the context of the intent it belongs to.

Hypothesis-driven. Every action recommendation includes a predicted outcome, a specific timeframe, and a measurable metric by which success will be evaluated. Vague recommendations like "improve performance" are not acceptable.

Competitive context matters. Never recommend bid reductions without checking the competitive landscape for that intent. A bid cut on a contested intent may cede share to an aggressive competitor who will be difficult to displace once entrenched.

Cost is truth. Use True Contribution Margin — including full landed cost — for all profitability calculations. Amazon's fee view is incomplete and systematically overstates your margin.

THE DECISION CHAIN IN PRACTICE

Phase 1: Snapshot

The agent begins every weekly session by pulling the complete picture:

> "Show me the current state of all my defined intents for the last 30 days. I want aggregated purchase share, click share, and impression share for each one, the opportunity classification, and which intents have moved significantly since last week."

> "Compare my top five intents side by side — show me which ones are improving, which are stable, and which are declining."

> "Show me the intent-by-ASIN matrix — which of my products are performing strongly within each intent, and where are the gaps?"

The result is a comprehensive intent performance snapshot: the raw material for all subsequent analysis.

Phase 2: Constraint Check

Before identifying any opportunities, the agent checks hard constraints:

> "Check my current inventory health. For every active SKU, tell me the days of supply based on this week's actual velocity, the throttle mode each product should be in, and flag any SKU where the mode has changed since last week."

> "Show me my True Contribution Margin and maximum break-even CPC for each active product, using the most recent landed cost data."

The constraint check is not a formality — it is the gate that all subsequent recommendations pass through. Any ASIN in Conservation or Emergency mode is flagged immediately and excluded from growth recommendations for that week.

Phase 3: Opportunity Scan

With constraints established, the agent applies the analytical frameworks:

> "Run the intent opportunity ranking for this week. Give me the top 15 opportunities ranked by priority score, classified by type — hidden gems, leaky intents, dominated intents, and untargeted intents. For each one, flag which ASINs are cleared for action based on inventory constraints."

> "Run the full intent gap analysis. Which high-volume intents have purchase share below my Yellow threshold and have not been targeted this week?"

> "Run the incrementality check across all intents. Which intents have incrementality scores above 0.8 — meaning PPC is almost entirely driving my presence? Which have scores below 0.2 — meaning I may be wasting ad spend?"

> "Check for funnel leaks. Are there any intents where my impression share is strong but click share or conversion share is significantly weaker than it was two weeks ago?"

Phase 4: Action Generation

From the opportunity scan, the agent generates a prioritized action list. Each recommendation follows the same structure:

> "Generate this week's intent actions, respecting all inventory constraints. For each recommendation, tell me the specific action, the data rationale behind it, the hypothesis including expected outcome and timeframe, the constraint status confirming which ASINs are cleared, and the priority score."

A well-formed recommendation looks like this:

Priority 1 — Increase bids on Pillow for Hot Sleepers intent, ASIN1 only.

Rationale: This intent has a conversion power score of 47% — meaning shoppers who do find your listing are buying at a strong rate — but your impression share is only 4%. This is a classic hidden gem: under-visible, high-converting. ASIN1 is in Aggressive mode with 67 days of supply. Max break-even CPC is $1.59.

Hypothesis: A bid increase from $0.85 to $1.05 will move impression share from 4% to 8% and purchase share from 2% to 4% within 14 days.

Constraint: Cleared. ASIN1 is in Aggressive mode.

Priority 2 — Launch A+ content update for Pillow for Neck Pain intent, all ASINs.

Rationale: This intent is classified as "leaky" — your impression share is 12% and click share is 8%, but purchase share is only 7%. You are visible, you are being clicked, but you are not closing. This is a persuasion problem, not a visibility problem. No bid increase will fix it.

Hypothesis: Updated A+ content addressing neck pain relief benefits specifically will improve the conversion rate by 2 percentage points within 30 days. Purchase share will rise from 7% to 9% without additional ad spend.

Constraint: No inventory impact. Proceed immediately.

Priority 3 — No action on Cooling Memory Foam intent this week.

Rationale: A competitor holds 42% purchase share on this intent. Your current share is 1.5%. The intent is contested by a dominant player. The cost of dislodging an entrenched competitor is high relative to the probability of success. Observe for four weeks to see if their share shows any vulnerability.

Constraint: Regardless of inventory status, this intent is not worth contesting at current budget levels.

Phase 5: Logging

Every session ends with documentation:

```
"Log today's full analysis — the snapshot data, constraint status for each ASIN,
the opportunities identified, and all actions recommended with their hypotheses.
Format it for my action log so I can review it during next week's governance
loop."
```

The log entry becomes the baseline for next week's review. When you open the governance loop next Monday, the first question is always: did what we predicted actually happen?

SECTION VII: PRACTICAL WORKSHOP

Chapter 13: Running the Intent Machine on Your Product

HANDS-ON DATA EXTRACTION AND FIRST INTENT ANALYSIS

This is where the book becomes real.

Everything you have read — the intent gap analysis, the incrementality framework, the blue ocean discovery, the marketing throttle, the governance loop — is now going to be applied to your product, with your data, in real time.

This chapter is a hands-on walkthrough. You will connect to your NeonPanel data, run a structured sequence of AI-guided analyses, and produce the complete intent dossier for your product.

PREREQUISITES

Before you begin, ensure:

1. **NeonPanel account is active** with your Amazon Seller Central integration connected and synchronized. Data should be current — synced within the last 24 hours.

1. **Brand Analytics access is enabled.** This requires Brand Registry. If you are not yet brand-registered, you can still use the advertising and financial tools, but the Brand Analytics frameworks — SQP, SCP, and intent analytics — will not be available.

1. **At least one Sponsored Products campaign is running.** You need PPC data to run the incrementality analysis. Even a single campaign with a few weeks of history is enough to start.

1. **Your target Parent ASIN is identified.** Pick the product you want to analyze. Ideally your best-seller or your highest-potential product — the one where a 5% increase in purchase share across a major intent would have the most revenue impact.

1. **Your AI environment is configured.** NeonPanel should be accessible via your preferred AI assistant, so you can run natural-language queries and get structured data responses.

THE EXTRACTION SEQUENCE

The workshop follows four phases. Each phase can stand alone after the setup phase completes, but they build on each other and are most powerful when run in sequence.

Phase 0: Bootstrap Your Intents

If this is your first time, you have no intent clusters yet. This phase creates them.

Start by pulling your top search terms from Brand Analytics:

> "Show me my top 100 search terms by search volume from the last 30 days. For each one, include the search volume, my brand's purchase share, and the current Red-Yellow-Green classification."

Save this list. It is the raw material for your intent taxonomy.

Now ask the AI to cluster them:

> "Here are my top 100 search terms. Please cluster them into 8 to 12 intent groups based on what customers are actually trying to achieve — not just keyword similarity, but underlying purchase motivation. For each proposed intent, give it a clear customer-facing name, describe the goal the customer is trying to fulfill, list which search terms belong to it, and tell me your confidence level in each grouping. Also flag any terms that are genuinely ambiguous and could belong to multiple intents."

Review the proposed clusters carefully. This is the most important 10 minutes in your setup. Ask clarifying questions. Refine clusters that do not match how your customers actually think.

Once you are satisfied:

> "Please save the following intent clusters to my account." (Then describe each one you want to keep, with any refinements you have made.)

Phase 1: Intent Performance Analysis

With your intents defined, pull the performance data:

> "Analyze my performance across all defined intents for the last 30 days. For each intent, show me the total search volume across all terms in the cluster, my brand's aggregated impression share, click share, and purchase share, the opportunity classification, and which individual search terms within each intent are driving the most conversion."

This gives you the complete map of your competitive position at the intent level — the foundation of every strategic decision.

Next, compare your intents to understand relative priority:

> "Show me a side-by-side comparison of my top five intents ranked by conversion share. Which ones am I winning on conversion? Which ones have strong visibility but weak conversion? Which ones am I barely present on at all?"

Then examine the ASIN-level picture:

> "Show me how each of my individual ASINs is performing within each intent. I want to see which products are strong in which customer segments and which products are consistently underperforming across all intents."

This cross-analysis often reveals that one product is carrying the weight across multiple intents while another is deadweight — a finding that has significant implications for both listing optimization and advertising strategy.

Phase 2: Opportunity Identification

With the performance landscape clear, identify where to act:

> "Run the full opportunity ranking across all my intents. Classify each intent by type — hidden gem (high conversion potential, low visibility), leaky (high visibility, low conversion), dominated (competitor with more than 40% share), or underdeveloped (I have some presence but no coordinated strategy). Rank them by priority and tell me the single most important action for each."

Drill deeper on your top gap intents:

> "For the Pillow for Neck Pain intent, show me the detailed funnel breakdown. Where exactly am I losing shoppers — at the impression stage, the click stage, or the conversion stage? Break it down by each of my ASINs so I can see if the leak is product-specific or category-wide."

Run the incrementality check:

> "For each of my active intents, calculate the incrementality score. Tell me which intents are Growth intents where PPC is my only visibility, which are Flywheel intents where organic is beginning to take over, and which are Waste intents where I am paying for clicks I would earn organically anyway."

Run the blue ocean scan:

> "Which of my intents have purchase share above 2% but zero PPC spend — meaning I have organic traction I have never amplified? And are there any major search terms in my category that do not appear in any of my current intent clusters — terms the market cares about that I am completely absent from?"

Phase 3: Operations and Constraints

Before acting on any opportunity, establish the financial and inventory boundaries:

> "Show me my full profitability picture per ASIN — True Contribution Margin after all costs including landed cost, the maximum break-even CPC I can afford for each product, and how FIFO batch costs are currently affecting my margins."

> "Check my current inventory health. For every active SKU, tell me the days of supply based on this week's actual velocity, the throttle mode each product is in, and which products have changed modes since my last check."

> "Given my inventory constraints, which intents and which ASINs are cleared for aggressive growth actions this week? Which are in Maintenance or Conservation mode and should be held steady?"

This constraint check turns the opportunity list into an actionable list — filtered for what you can actually do right now versus what you need to wait on.

YOUR INTENT DOSSIER

When the extraction is complete, you have five deliverables:

1. Intent Landscape Map. All your defined intents with their total market size, your current purchase share, RYG classification, and momentum direction. This is the strategic overview of your competitive position.

2. Opportunity Ranking. Intents ranked by priority type — hidden gems, leaky intents, dominated intents, and untargeted intents — with the single most important recommended action for each.

3. ASIN by Intent Matrix. A cross-tabulation showing which of your products perform strongly within each intent and where the gaps are. This identifies listing optimization priorities and product-specific advertising strategies.

4. Constrained Action List. This week's recommended actions per intent, with data rationale, hypothesis, expected outcome, and confirmation that each action respects your current inventory and financial constraints.

5. Week 1 Governance Baseline. A timestamped snapshot of purchase share, click share, conversion metrics, and inventory status — the starting point against which all future weekly deltas will be measured.

With these five deliverables in hand, you are no longer managing keywords. You are managing intents. Next Monday, you run the loop again. You compare Week 2 to Week 1. You evaluate your hypotheses. You adjust. You learn.

Welcome to systematic growth.

Chapter 14: From Keyword to Intent — A Practical Example

CASE STUDY: HOW A PILLOW SELLER DOUBLED REVENUE USING INTENT-CENTRIC STRATEGY

The Seller: Mid-market private-label pillow brand, 4 ASINs, $180K/month revenue, $35K/month ad spend.

The Problem: After running keyword-level growth for 18 months, the seller had plateaued. They were bidding on 200+ keywords, their ACoS was 19% (acceptable), but their market share had stalled. They had no idea whether they were winning or losing.

PHASE 1: BOOTSTRAP INTENTS

The seller exported their top 120 SQP keywords. Used LLM clustering to propose intents. The agent suggested 9 intents:

1. Pillow for Neck Pain (27 keywords)
2. Pillow for Side Sleepers (19 keywords)
3. Cooling Pillow (16 keywords)
4. Pillow for Hot Sleepers (14 keywords)
5. Pillow for Allergies (12 keywords)
6. Premium/Luxury Pillow (11 keywords)
7. Budget Pillow (8 keywords)
8. Pillow for Back Sleepers (9 keywords)
9. Bamboo/Eco Pillow (4 keywords)

The seller reviewed the clusters, refined a few (merged Bamboo/Eco into Premium because they were both niche/high-price), and persisted 8 intents.

PHASE 2: INTENT ANALYSIS

Using `brand_analytics_analyze_intent_performance`, the seller discovered:

Intent	Size	Share	Classification	Spend	ROAS
Pillow for Neck Pain	247K	7%	Leaky	$8,200/mo	1.8x
Pillow for Side Sleepers	185K	11%	Dominant	$6,500/mo	2.1x
Cooling Pillow	92K	3%	Red	$4,200/mo	1.2x
Pillow for Hot Sleepers	68K	4%	Hidden Gem	$1,200/mo	3.2x
Pillow for Allergies	54K	8%	Yellow	$3,100/mo	2.0x
Premium/Luxury	41K	5%	Red	$2,800/mo	1.5x
Budget Pillow	156K	9%	Yellow	$5,200/mo	1.6x
Back Sleepers	62K	6%	Yellow	$3,800/mo	1.9x

Insights:

1. **Pillow for Hot Sleepers is a hidden gem.** Only 68K searches, but 4% share and 3.2x ROAS. This intent is underinvested.

1. **Pillow for Neck Pain is leaky.** 247K searches (the largest), but only 7% share despite high spend. Buyers are clicking but not converting. Likely a listing quality issue.

1. **Cooling Pillow is a drain.** Low share (3%), declining searches, 1.2x ROAS. Not worth the spend.

1. **Budget Pillow is undervalued.** 156K searches, decent share (9%), 1.6x ROAS. Could be grown with better targeting.

PHASE 3: ACTION PLAN

Week 1–2: Reallocate Budget

- Pause Cooling Pillow ads. Redirect $4.2K/mo to Pillow for Hot Sleepers.
- Increase bid on Pillow for Hot Sleepers by $0.25 per keyword.
- Reduce bids on Pillow for Neck Pain to conserve spend, but increase budget on the hidden gem.

Hypothesis: Hidden gem ROAS (3.2x) is higher than neck pain (1.8x). Shifting $4.2K from drain to gem will improve overall ROAS by 0.4x (from 1.87x to 2.27x).

Week 3–4: Address Neck Pain Leakiness

- Analyze funnel. Are clicks high but conversions low? Yes (12% click share, 7% conversion share).
- Update main product image to emphasize neck pain relief. Test new A+ content module.
- Maintain PPC spend; focus on fixing the listing, not bidding harder.

Hypothesis: Improved listing will increase conversion rate by 2 percentage points. Purchase share will rise from 7% to 10% in 30 days despite flat spend.

Week 5–8: Expand Budget Pillow

- Budget Pillow has 156K searches and decent 9% share. Increase bids by $0.15.
- Add new keyword variants targeting "budget pillow under $30" and "affordable pillow."
- Test if price emphasis resonates with this price-conscious segment.

Hypothesis: Price-focused messaging will improve click share and conversion share by 2% and 1.5% respectively.

RESULTS (8 WEEKS LATER)

- **Pillow for Hot Sleepers:** Share rose from 4% to 7% (+75% improvement). ROAS remained 3.2x.
- **Pillow for Neck Pain:** Share rose from 7% to 9% (+29%) after listing update. Conversion rate improved 1.8 percentage points.
- **Budget Pillow:** Share rose from 9% to 11% (+22%) after price-focused keyword expansion.
- **Overall:** Ad spend held flat at $35K/mo, but revenue rose from $180K to $245K (+36%).

The seller had not invented new products or dramatically changed their positioning. They had simply reorganized their thinking from keywords to intents, reallocated budget to higher-opportunity intents, and fixed underlying funnel problems.

Intent-centric analysis made the pattern visible. Keyword-level analysis never would have.

Chapter 15: Integration with External Data

USING BRIGHT DATA AND SUPABASE TO TRIANGULATE INTENT PERFORMANCE

Brand Analytics tells you *what* is happening — impressions, clicks, conversions, purchase share. It is precise and trustworthy. But it does not tell you *why* it is happening.

Your conversion share on the Pillow for Neck Pain intent dropped 2% this week. Is it because a competitor launched a 20% coupon? Because they received a Best Seller badge? Because their newest review mentions neck pain specifically and is pulling shoppers who were previously comparing you favorably? Brand Analytics cannot tell you. You have the outcome; you do not have the cause.

External data fills this gap.

AUGMENTING INTENT DATA WITH SERP SNAPSHOTS

By integrating a web scraping tool such as Bright Data alongside your NeonPanel analytics, your AI agent can pull Amazon search result snapshots for representative keywords within each intent. The idea is to combine the what from Brand Analytics with the why from real-time SERP observations.

Ask the AI:

> "Search Amazon for 'orthopedic neck pillow' — the primary keyword in the Pillow for Neck Pain intent — and show me the top 10 results. For each result, I want to see the price, star rating, total review count, and any badges displayed — Best Seller, Amazon's Choice, Climate Pledge Friendly, or similar."

The result might reveal: your product is priced 23% higher than the current leader, who also holds 12,400 reviews compared to your 3,200. You lack any badge. The leader recently earned Amazon's Choice status on this exact keyword.

These are not metrics Brand Analytics surfaces. But they explain exactly why your conversion share is declining despite stable impression share. Shoppers are seeing your product, considering it, and choosing the better-badged, better-reviewed, lower-priced alternative.

This is the qualitative context that turns a Brand Analytics diagnosis into an actionable response.

BUILDING A COMPETITIVE TIMELINE WITH SUPABASE

Spotting a single week's SERP snapshot is useful. Building a historical timeline of how the competitive landscape has evolved is transformative.

By persisting weekly SERP snapshots in a structured database such as Supabase, you create a timeline that can be correlated against your Brand Analytics data. Over time, a clear causal picture emerges.

For example, a weekly log might show:

Week 1: The top competitor for the Pillow for Neck Pain intent is priced at $39.99 with 12,400 reviews. Your purchase share on the intent is 9%.

Week 2: The same competitor dropped their price by $5 to $34.99. Their review count grew to 12,600. Your purchase share dropped to 7%.

Week 3: Their price returned to $39.99. Your purchase share partially recovered to 8%.

This three-week pattern tells you something important: your brand has price-sensitive buyers on this intent. When the leader is cheaper by more than 10%, you lose share. When pricing is comparable, you hold your position. That insight is not in Brand Analytics — it emerges from the intersection of market observation and performance data.

Ask the AI:

> "Based on the last 8 weeks of SERP snapshots for the Pillow for Neck Pain intent, are there any patterns between competitor pricing changes, badge changes, or review velocity and my own purchase share movement? I am looking for causal relationships I can use to make proactive decisions rather than reactive ones."

The goal is to move from a seller who reacts to lost share to a seller who anticipates competitive moves and responds before the damage compounds.

Afterword

Amazon's marketplace rewards three things: relevance, velocity, and consistency.

Relevance means appearing in front of customers expressing a specific intent. Velocity means converting those customers faster than alternatives. Consistency means doing this, week after week, compounding your advantages.

Intent-centric analytics make all three possible.

The keyword era is closing. The intent era is opening. The sellers who make the shift will have a 24–36 month advantage before the rest of the market catches up.

You have the frameworks. You have the tools. You have NeonPanel's intent analytics.

Build the machine. Run the loop. Grow.

APPENDIX: QUICK REFERENCE — NEONPANEL INTENT TOOLS

Tier 1: Intent Lifecycle

Create Intent Cluster Manually

Define a new intent by giving it a name, a description of the customer goal it represents, the list of search terms that belong to it, and your confidence in the grouping. The tool saves the cluster and returns a unique intent ID you can reference in all subsequent analyses.

AI-Assisted Intent Clustering

Provide a list of your top search terms and specify how many intent clusters to propose — typically 8 to 12. The AI analyzes the terms for underlying customer intent, proposes clusters with names, descriptions, mapped terms, and confidence scores, flags ambiguous terms that could reasonably belong to more than one intent, and recommends any merges or splits worth considering.

Map Search Terms to Intents

Assign specific search terms to an existing intent. Supports one-to-many mapping, so a term that genuinely belongs to two intents can be assigned to both. Flags conflicts where the same term has been assigned to intents with very different customer goals.

List and Browse Intent Clusters

View all your defined intents with filtering and sorting options. Includes summary statistics for each cluster — how many terms it contains, its current performance classification, and when it was last updated.

Tier 2: Intent Analytics

Analyze Intent Performance

The primary analytical tool. For a given intent, aggregates all Brand Analytics metrics — impressions, clicks, cart adds, purchases — across every search term in the cluster. Returns your

brand's aggregated impression share, click share, and purchase share, a composite visibility index and conversion power score, an opportunity classification (hidden gem, leaky, dominated, or competitive), performance trends over the specified time window, and the top-performing search terms and ASINs within the intent.

Intent vs. Intent Comparison

Compare two to five intents side by side on a specified metric — purchase share, conversion rate, impression share, or ROAS. Returns a ranked comparison matrix, highlights each intent's relative strengths and weaknesses, and provides a recommendation on which to prioritize for the current period.

Intent by ASIN Matrix

Produces a pivot table showing how each of your ASINs performs within each intent, for the metric you specify. Identifies hotspots (ASIN and intent combinations where performance is unusually strong) and cold spots (combinations where you have inventory and organic presence but poor conversion). Used to determine which products to prioritize in each intent's advertising campaigns.

Intent Opportunity Ranking

Ranks all your defined intents by opportunity score, filtered by opportunity type: hidden gems (strong conversion potential, low visibility), leaky intents (high visibility but poor conversion), dominated intents (a competitor holds more than 40% purchase share), or untargeted intents (you have organic presence but no ad spend). Returns a recommended action for each, with an estimated revenue uplift if the recommended action is executed.

Tier 3: Intent Playbooks

Intent Growth Machine Diagnosis

Runs the full diagnostic framework across all your intents — gap analysis, incrementality check, funnel diagnosis, and throttle constraint check — in a single pass. Returns a prioritized list of recommended actions with hypotheses, constraint status per ASIN, and confidence scores based on your historical hypothesis accuracy. This is the tool that powers the weekly NeonaSphera Agent session.

Write and Update Intent Clusters

Create, update, rename, or deactivate intent clusters. Supports bulk operations for updating multiple intents or reassigning search terms. When an intent is deactivated, its historical performance data is preserved so you can reactivate and continue trending analysis.

End of Book