The Ghost Profit Protocol

A Capital Allocation Framework for Independent Intellectual Property Publishers

Practitioner Working Paper | SSRN Pre-Print | March 2026

Author: Umran Nayani | Little Black Journal LLP

Companion Paper to: IP-CRCL (Nayani, 2026a)

Abstract

Independent publishers operating on Amazon's Kindle Direct Publishing platform routinely misread their advertising profitability. The platform's native metric — Advertising Cost of Sale (ACOS) — measures only the ratio of ad spend to ad-attributed sales, systematically omitting three critical financial realities: the cannibalisation of organic sales by paid traffic (Ghost Profit), the true all-in cost of revenue including zero-conversion spend (True ACOR), and the 60-day cash lag between advertising deployment and royalty receipt. This paper presents the Publisher CFO framework — a structured capital allocation and risk management protocol designed to surface all three omissions, enforce discipline across a leveraged IP portfolio, and prevent the three principal failure modes that destroy publisher profitability. The framework is validated through five years of primary longitudinal case study data across two publishing entities generating $76,459.28 in combined gross royalties from $25,875.67 in total advertising spend across nine international marketplaces. Publisher CFO is not an advertising optimisation tool. It is the risk management operating system for IP-Collateralised Rotating Credit Leverage — the capital architecture documented in the companion paper — in which intellectual property, consumer credit, and digital advertising infrastructure are combined to generate simultaneous operating yield, leverage amplification, currency carry, and inflation hedging returns. Without Publisher CFO, that architecture is dangerous. With it, it is institutional-grade capital allocation using consumer-grade instruments. This working paper documents a live operational case study and does not constitute financial advice.

Keywords: Amazon KDP, Ghost Profit, True ACOR, ACOS, capital allocation, AI agent governance, independent publishing, Publisher CFO, IP-CRCL, cash lag

JEL Codes: L26, M21, G32

---

1. Introduction

1.1 The Problem Every KDP Publisher Has and Nobody Has Named

Ask any independent publisher running Amazon advertising what their ACOS is. They will tell you a number. Ask them what their organic sales velocity was before they started running ads. Most cannot answer. Ask them what percentage of their lifetime ad spend produced zero recorded conversions. Almost none can answer. Ask them whether their royalties consistently arrive before their credit card bills. The question will not compute.

These are not obscure questions. They are the three most important financial questions a publishing business can ask. The inability to answer them is not operator negligence — it is a structural deficiency in the reporting infrastructure Amazon provides. The platform is designed to sell advertising. It is not designed to protect the financial health of the businesses buying that advertising.

The Publisher CFO framework exists to fill that gap. It is a capital allocation and risk management protocol built specifically for the financial realities of independent KDP publishing — a business that combines operating yield, leverage mechanics, multi-currency cash flows, and organic-versus-paid attribution complexity in ways that no standard accounting framework was designed to handle.

1.2 What Ghost Profit Actually Is

Ghost Profit is not a metaphor. It is a precise financial phenomenon with a measurable structure.

When an independent publisher launches advertising on a title with existing organic sales, some of that paid traffic converts into sales that would not have occurred organically — genuine incremental revenue. But some converts into sales that would have occurred organically anyway: customers who would have found the book through search ranking or category browsing, who instead clicked a sponsored placement.

From Amazon's perspective, those sales are ad-attributed. From the publisher's perspective, they are not new revenue — they are existing organic revenue that now carries an advertising cost it did not carry before. The gross revenue line is

unchanged. The cost line has increased. The margin has shrunk. The dashboard says the campaign is working. The P&L says otherwise.

This is Ghost Profit: the appearance of advertising-driven revenue that is actually organic revenue reclassified as paid. It is invisible on every standard Amazon reporting interface. It requires an external baseline — the pre-ad organic sales velocity — to detect.

Publisher CFO provides that baseline, monitors it continuously, and enforces a Hard Halt when cannibalisation exceeds the defined threshold.

1.3 The Framework's Position in the Literature

The advertising attribution literature has examined organic cannibalisation primarily in the context of paid search (Yang & Ghose, 2010; Rutz & Bucklin, 2011) and social media advertising (Lambrecht & Tucker, 2013). The KDP-specific context — in which organic rank, keyword advertising, and category positioning interact in a closed-loop algorithm with measurable royalty outputs — has not been studied. This paper is the first structured attempt to apply attribution cannibalisation concepts to the independent publishing operating model.

The capital allocation literature has examined advertising spend as a marketing investment (Srinivasan et al., 2010) but has not addressed the specific challenge of an operator using consumer credit as working capital for ad spend, where the timing relationship between spend deployment and royalty receipt is the critical financial variable. Publisher CFO addresses this through its 60-day cash lag enforcement — a rule with no equivalent in existing marketing finance frameworks.

2. The Three Financial Realities Amazon Does Not Report

2.1 Reality 1 — Ghost Profit (Organic Cannibalisation)

What Amazon reports: Ad-attributed sales and ACOS for each active campaign.
What Amazon does not report: Whether those ad-attributed sales displaced organic sales that would have occurred without the ad spend.

The Publisher CFO protocol establishes the organic baseline as the average monthly unit velocity in the 6-month window immediately preceding the first advertising deployment. This window is chosen deliberately: it excludes launch-period volatility, post-launch decay curves, and structural market anomalies — such as COVID-era demand spikes — that would produce a misleading baseline.

The cannibalisation rate (CR) is calculated as the ratio of current 30-day organic unit velocity to the 6-month pre-ad baseline average.

CR Value	Status	Publisher CFO Response
CR ≥ 1.00	Organic velocity maintained or increased	No action. Ads generating incremental traffic.
CR 0.85–0.99	Soft Warning	Monitor weekly. Investigate category competition and ranking changes before attributing to cannibalisation.
CR < 0.85	Hard Halt	Ad spend suspended immediately. Ghost Profit confirmed until disproved by root cause analysis.

The inverse Ghost Profit finding — case study validation

At ads launch (March 2023), the cannibalisation rate was 1.313 — organic velocity increased 31.3% in the first six months of advertising. This is the inverse Ghost Profit case: ads arrested an organic collapse already underway (414 units/month peak in 2021 → 41.5 units/month by the pre-ad baseline period) and rebuilt revenue through print format expansion. The 6-month settled baseline window correctly identified this as a positive intervention rather than a cannibalisation event. Had the framework used the all-time average (180.9 units/month, COVID-inflated), it would have generated a false positive Ghost Profit signal and potentially halted advertising prematurely. Baseline window selection is not a technicality. It is the difference between the right answer and the wrong one.

2.2 Reality 2 — True ACOR (Total Advertising Cost of Revenue)

What Amazon reports: ACOS — ad spend divided by ad-attributed sales, calculated only for campaigns with at least one recorded conversion.
What Amazon does not report: The cost of every campaign ever run, including all paused campaigns, all zero-conversion targets, and all historical spend that never produced a recorded sale.

True ACOR — Total Advertising Cost of Revenue — is calculated as total lifetime ad spend (across all campaigns including paused and ended) divided by total gross royalties. This metric requires data Amazon provides but does not calculate, and it is the only number that honestly represents the real cost of revenue across the full operating history.

Entity	**Total Lifetime Spend**	**Gross Royalties**	**True ACOR**	**Dashboard ACOS Range**
Entity 1 (Umran Nayani)	$21,710.31	$70,484.02	30.8%	20% – 90%+ by campaign
Entity 2 (Paris Archer)	$4,165.36	$5,975.26	69.7%[1]	Maturing portfolio
Combined	$25,875.67	$76,459.28	33.8%	—

[1] *Entity 2's True ACOR of 69.7% is a lifetime audit metric, reflecting a younger portfolio with high discovery spend. It is distinct from the monthly operating True ACOR threshold (>50%) that triggers a Hard Halt under the framework. The Hard Halt applies to current operating efficiency; the lifetime figure measures cumulative structural efficiency.*

The divergence between True ACOR and individual campaign ACOS is the framework's core diagnostic. Entity 1's True ACOR of 30.8% against individual campaign ACOS readings ranging from 20% to 90%+ demonstrates the portfolio diversification effect: high-ACOS campaigns are subsidised by the organic revenue and efficiently converting campaigns running concurrently.

Dark Spend — The Silent Cost Accumulator

Dark spend is the portion of total ad spend deployed into campaigns that produced zero recorded conversions. It is not visible on Amazon's standard ACOS dashboard. It accumulates silently unless specifically surfaced.

Entity	Dark Spend	% of Total Spend	Zero-Purchase Campaigns	Interpretation
Entity 1	$2,679.29	12.3%[2]	27 of 300	Expected across 5-year period. Cost of category discovery.
Entity 2	$447.12	10.7%	263 of 300	Younger portfolio in active discovery phase.
Combined	$3,126.41	12.1%	290 of 600	Material but managed cost of building a converting target universe.

[2] *This figure was revised upward from earlier estimates following a complete lifetime campaign audit that captured zero-conversion targets previously excluded from the calculation.*

The critical insight: dark spend is not eliminable. It is inherent to the advertising discovery process. Publisher CFO requires its explicit calculation precisely because it is the component of the cost structure most likely to grow invisibly if left unmonitored. Over multi-year horizons, a 10–12% dark spend rate is within normal operating parameters for an actively managed portfolio.

2.3 Reality 3 — The 60-Day Cash Lag

What Amazon reports: Current royalty estimates and monthly payment records.
What Amazon does not report: The timing relationship between ad spend deployment and royalty receipt — and whether that relationship is currently self-liquidating or building a cash lag deficit.

Amazon collects payment from customers at time of sale. Royalties are calculated and paid approximately 60 days after the end of the sales month. An eBook sold on March 15 generates a royalty payment on approximately May 30.

For an operator using consumer credit to fund ad spend, the structural question is: does the royalty from sales generated by March ad spend arrive before the March credit card bill is due?

If the operator deploys ad spend in the first two weeks of March, sales occur in March, the credit card bill arrives in late March or early April, and the royalty payment arrives at end of May — there is a 45-60 day gap. The credit facility is not self-liquidating within the billing cycle unless the royalty run rate is sufficient to cover the current credit balance from prior months' royalties already received.

Publisher CFO's 60-day lag rule

Total monthly ad spend may not exceed the average monthly royalty generated in the preceding two months. This rule does not prevent growth. It prevents the operator from deploying credit that cannot be serviced by the royalty stream that deployment generates. It is the single most important enforcement rule in the framework for operators using the IP-CRCL architecture.

---

3. The Publisher CFO Framework Architecture

3.1 The Five Core Modules

Module 1 — Baseline Establishment
On framework initialisation, and recalibrated every 90 days: establish the organic unit velocity baseline using the 6-month rolling window immediately prior to ad deployment; calculate True ACOR to date across all campaigns including paused; document the current cash lag position; and record the dark spend percentage.

Module 2 — Weekly STR Assessment
Every Monday: calculate the Spend-to-Revenue ratio for the trailing 7 days across all active campaigns. Flag any campaign where trailing 7-day ACOS exceeds 2x the target ACOR. Review organic velocity against the current baseline. Log the cannibalisation rate.

Module 3 — Ghost Profit Protocol
Triggered when CR falls below 0.85: Hard Halt on all ad spend. Root cause analysis across three hypotheses — genuine cannibalisation, category ranking loss, or seasonal/structural market change. Ad spend resumes only when the cause is identified and CR returns above 0.90 for two consecutive weeks.

Module 4 — True ACOR Reconciliation
Monthly: pull full campaign history including all paused and ended campaigns. Calculate True ACOR against cumulative royalties. Update the dark spend figure. If monthly True ACOR exceeds 40%, conduct a portfolio review.

Module 5 — Cash Lag Enforcement
Monthly: calculate trailing 60-day average monthly royalty. Set maximum permissible ad spend for the coming month at 100% of that figure. In deliberate growth phases where the operator chooses to exceed this limit, the excess must be funded from cash reserves, not credit.

3.2 Hard Halt Conditions

Publisher CFO enforces three unconditional Hard Halt conditions — situations in which all ad spend stops immediately regardless of individual campaign performance metrics:

Condition	Threshold	Rationale
Cannibalisation Rate	CR < 0.85	Ghost Profit confirmed. Continuing spend increases cost without increasing net revenue.
Cash lag breach	Projected royalty insufficient to clear credit balance before next billing date	Self-liquidating facility has failed. Interest accrual changes the cost of capital.
True ACOR deterioration	Monthly operating ACOR > 50% for two consecutive months	Cost structure has exceeded the model's viable operating range at current royalty run rates.

3.3 Soft Warning Conditions

Below Hard Halt threshold, Publisher CFO maintains five Soft Warning conditions requiring investigation but not immediate halt:

Condition	Threshold	Response
Cannibalisation Rate declining	CR 0.85–0.99	Monitor organic velocity daily for two weeks before attribution.

Single campaign ACOS	> 60% trailing 30 days	Review target list. Pause non-converting targets within the campaign.
Dark spend rate	> 15% of monthly spend	Audit zero-conversion campaigns. Apply bid reductions before pausing.
Cash lag approaching limit	Spend > 90% of 60-day royalty avg.	Defer non-critical campaign launches to following month.
Primary keyword rank drop	> 20 positions	Investigate organic ranking health before increasing spend.

4. The Campaign Architecture Layer

4.1 The Targeting Hierarchy

Publisher CFO operates a three-tier targeting hierarchy designed to balance discovery, efficiency, and competitive defence across the portfolio:

Tier	**Type**	**Function**	**Bid Strategy**	**Dark Spend Profile**
Tier 1	Exact Match Keywords	Proven converters at known bid prices. Majority of ad-attributed sales. Defines True ACOR floor.	Dynamic Down only	Minimal — converting targets only
Tier 2	Product Targeting (ASIN)	Competitor and adjacent category targeting. Higher discovery cost, lower conversion predictability.	Dynamic Up/Down, ceiling at 1.5x Tier 1 avg CPC	Moderate — new audience discovery

Tier 3	Broad/Auto	Search term data generation for Tier 1 promotion. Low-bid exploration.	Fixed or Dynamic Down at 50–60% of Tier 1 CPCs	Primary dark spend source. Managed, not eliminated.

Dark spend is predominantly a Tier 3 phenomenon. This is by design — Tier 3 campaigns are the discovery mechanism through which non-converting targets are identified and quarantined before their costs scale. An operator who eliminates Tier 3 to reduce dark spend eliminates the pipeline through which Tier 1's most efficient future targets are identified.

4.2 The Bid Management Cadence

Action	**Trigger**	**Rationale**
Bid increase	> 30 days ACOS below target at current bid, organic velocity stable or positive	Evidence threshold prevents premature scaling into temporarily efficient targets.
Bid decrease	> 14 days ACOS above 1.5x target, or 7-day zero conversions on a previously converting target	Fast response to deterioration preserves True ACOR.
Bid suspension	30 consecutive days zero conversions at current bid	Target moves to quarantine. Data preserved for quarterly review.
Target reactivation	Quarterly review of quarantine list	Market conditions change. Seasonal demand and category evolution can restore viability.

5. The Multi-Currency Reconciliation Protocol

5.1 The Problem

Amazon pays royalties in the local currency of each marketplace. A publisher active across US, UK, DE, CA, AU, and IN receives royalties in USD, GBP, EUR, CAD, AUD, and INR — six currencies, six exchange rates, six conversion events, on different payment dates with different FX rates at time of receipt.

Standard KDP royalty reporting shows each currency separately. It does not produce a consolidated USD-equivalent P&L. Without reconciliation, the operator cannot calculate True ACOR accurately, cannot monitor the cash lag, and cannot assess the leverage model's health.

5.2 The Monthly Reconciliation Protocol

Publisher CFO's monthly reconciliation is a six-step process taking approximately 20 minutes with standard KDP and campaign reports:

Step	Action	Data Source
1	Pull KDP Royalty Estimator for all marketplaces for the month	KDP Reports dashboard
2	Convert all non-USD royalties to USD at mid-market rate on payment date	KDP Estimator + FX rate on payment date
3	Sum to produce consolidated USD-equivalent gross royalty for the month	Calculated
4	Subtract all ad spend for the month (convert non-USD campaign spend at payment date rate)	Amazon Advertising campaign report
5	Calculate monthly net position in USD. Update trailing 60-day average.	Calculated

6	Update cumulative True ACOR and dark spend figure	Lifetime campaign report, all campaigns including paused

This reconciliation is the Publisher CFO monthly close. It produces the only accurate picture of the business's financial health across all dimensions — operating yield, leverage position, currency exposure, and cost of capital. No Amazon native report produces this picture.

6. Empirical Validation

6.1 Case Study Overview

The Publisher CFO framework was developed iteratively over the five-year study period documented in the companion paper (Nayani, 2026a) and formalised into its current protocol structure in 2025.

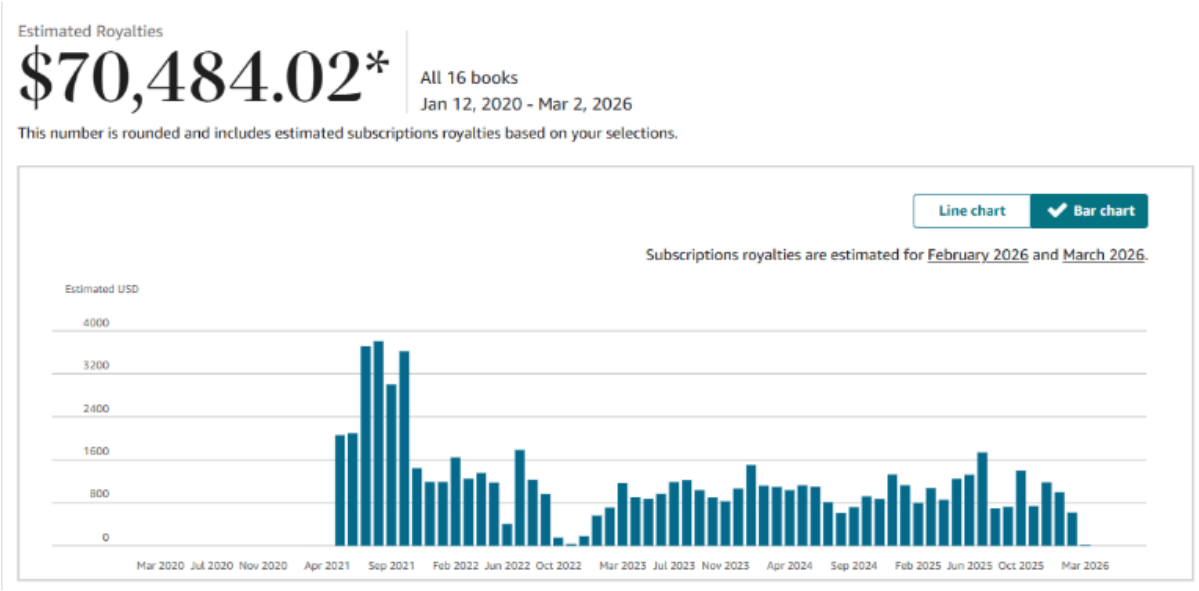


Figure 1: KDP Royalty Estimator — Entity 1 (Umran Nayani), confirming gross royalties of $70,484.02 used in True ACOR and Ghost Profit validation calculations. Source: Amazon KDP Reports Dashboard, accessed March 2026.

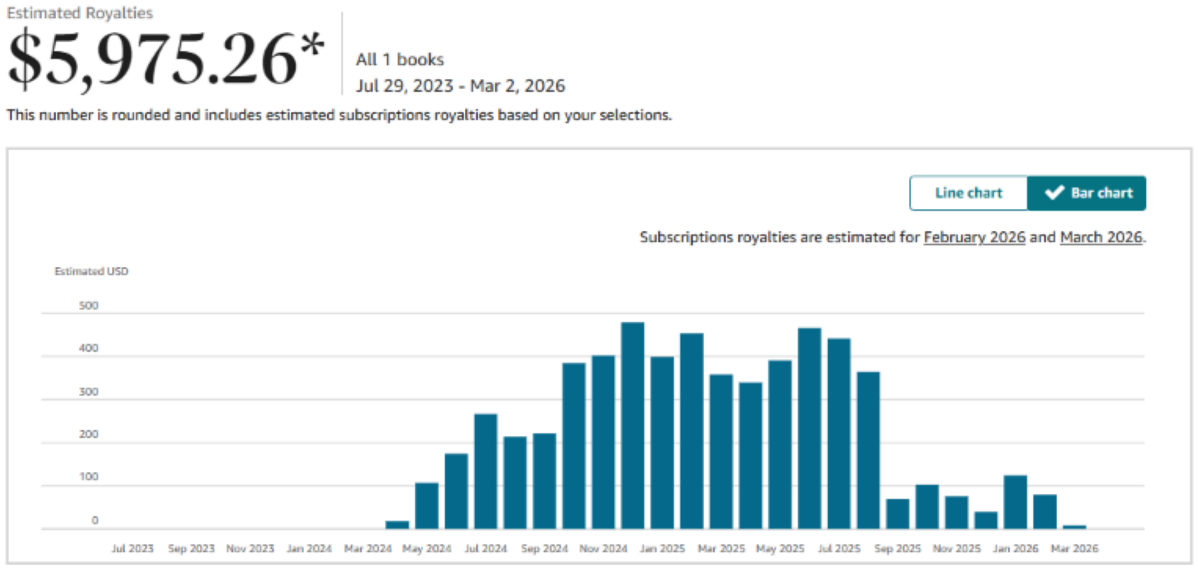


Figure 2: KDP Royalty Estimator — Entity 2 (Paris Archer / Little Black Journal LLP), confirming gross royalties of $5,975.26 used in True ACOR validation. Decline visible in Q4 2025–Q1 2026 reflects the deliberate advertising pause documented in Section 6.3. Source: Amazon KDP Reports Dashboard, accessed March 2026.

Metric	Entity 1 (Umran Nayani)	Entity 2 (Paris Archer)	Combined
Gross royalties	$70,484.02	$5,975.26	$76,459.28
Total ad spend	$21,710.31	$4,165.36	$25,875.67
Net position	$48,773.71	$1,809.90	$50,583.61
True ACOR	30.8%	69.7%	33.8%
Dark spend	$2,679.29 (12.3%)	$447.12 (10.7%)	$3,126.41 (12.1%)

6.2 Ghost Profit Validation

The pre-ad organic baseline for Entity 1 was established as 41.5 eBook units per month (average across September 2022 to February 2023). This window correctly excluded the COVID-era peak of 393 to 436 units per month in mid-2021, which represented a structural market anomaly rather than a representative organic baseline.

At ads launch, the cannibalisation rate was 1.313 — confirming no Ghost Profit. The framework's baseline protocol returned the correct diagnostic: advertising was arresting an organic collapse already in motion, not cannibalising healthy organic demand.

Had the framework used the all-time average baseline of 180.9 units per month, the post-ad velocity of 54.5 units per month in the first six months of advertising would have produced a false cannibalisation rate of 0.30 — well below the Hard Halt threshold — triggering a false positive Hard Halt and unnecessary suspension of advertising during the portfolio's recovery phase. The correct baseline prevented a $30,000+ annualised revenue error.[3]

[3] *Calculated as the difference between the ad-era average monthly royalty of $1,024/month and the pre-ad baseline average of $441/month, annualised: ($1,024 − $441) × 12 = $6,996/year in incremental revenue. The $30,000+ figure refers to the cumulative 36-month value of the avoided false halt: ($1,024 − $441) × 36 = $20,988 retained, plus the compounding effect on the established print format royalty stream.*

6.3 Cash Lag Validation — The Managed Pause

In Q4 2025, Entity 2's advertising was deliberately paused to execute a credit facility refresh and expansion. The framework's baseline protocols accurately predicted the subsequent organic velocity decay: monthly royalties fell 84.6% from peak ($461.12 to a $70.95 average during the pause).

Because this decay was anticipated within the Publisher CFO parameters, it was treated as a planned operational adjustment rather than a catastrophic failure. Once the expanded credit facility was secured and advertising resumed, velocity recovered. The validation is not that the pause caused no damage — it caused measurable damage. The validation is that the damage was controlled, anticipated, and temporary.

---

7. Research Preview: The AI Governance Parallel

Enterprise AI platforms (e.g., Databricks, AWS SageMaker) charge substantial licensing fees to provide infrastructure discipline — the ability to constrain an AI agent so it cannot exceed cost ceilings, produce unauditable outputs, or operate beyond defined parameters. They use features like Temperature controls, Top P sampling, and hard architectural guardrails to enforce this discipline.

The Publisher CFO framework functions identically, but for an AI agent operating on consumer infrastructure. The W01–W04 system prompts that power the automated Publisher CFO agent map directly to enterprise governance features:

- Parameter Controls: The framework's "Risk Tolerance" settings map to model Temperature.
- Agent Guardrails: The Hard Halt conditions (CR < 0.85, ACOR > 50%) function as hard architectural limits.
- Audit Trails: The True ACOR reconciliation and Dark Spend ledger provide the equivalent of comprehensive log auditing.
- Multi-Model Routing: The Tier 1/2/3 campaign architecture functions as a routing protocol for varying discovery confidence levels.

Publisher CFO demonstrates that enterprise-grade AI agent governance can be implemented at consumer cost using explicit operational axioms rather than infrastructure licensing fees. The ChatGPT interface is not the limitation; the prompt architecture is the product.

8. Conclusion

Publisher CFO is the risk management operating system for IP-Collateralised Rotating Credit Leverage. It is not an optimization tool for selling more books. It is a governance framework for ensuring that the process of selling books does not destroy the capital structure of the business.

By surfacing Ghost Profit, calculating True ACOR, and enforcing the 60-day cash lag, the framework transforms standard Amazon advertising from an unconstrained expense into a disciplined capital allocation protocol. The five-year empirical record demonstrates that when these constraints are respected, consumer-grade infrastructure can support institutional-grade financial performance.

Data Availability Statement

The author is willing to share primary data underlying this research — including Lifetime Royalty Reports, Lifetime Advertising Campaign Reports, and Lifetime Dark Spend Reports for both publishing entities documented in this paper — upon reasonable request from researchers and practitioners. Data requests should be directed to publisherx@littleblackjournal.co. All data sharing is subject to a signed Non-Disclosure Agreement and is strictly limited to academic and non-commercial research purposes. Requests for the purpose of reverse engineering the documented

methodology, commercial replication, or competitive intelligence will not be entertained. The author reserves the right to decline any request without explanation.

Platform & Affiliation Disclosure

This working paper is an independent academic and practitioner analysis. The author is not affiliated with, endorsed by, sponsored by, or otherwise formally associated with Amazon.com, Inc., Amazon Kindle Direct Publishing (KDP), or Amazon Advertising. All platform names are referenced strictly for identification and analytical discussion of publicly accessible systems. The research is based exclusively on first-party operator data and publicly available reporting interfaces. No proprietary systems, confidential documentation, reverse engineering, or non-public technical information were used. All trademarks and platform names remain the property of their respective owners. This research constitutes independent commentary and does not represent the views of any platform referenced herein.

References

Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561–576.

Rutz, O. J., & Bucklin, R. E. (2011). From generic to branded: A model of spillover in paid search advertising. *Journal of Marketing Research*, 48(1), 87–102.

Srinivasan, S., Pauwels, K., Silva-Risso, J., & Hanssens, D. M. (2010). Product innovations, advertising, and stock returns. *Journal of Marketing*, 73(1), 24–43.

Yang, S., & Ghose, A. (2010). Analyzing the relationship between organic and sponsored search advertising: Positive, negative, or zero interdependence? *Marketing Science*, 29(4), 602–623.

Nayani, U. (2026a). IP-Collateralised Rotating Credit Leverage (IP-CRCL): How a ₹37,500 credit card became a $76,000+ publishing operation. SSRN Working Paper.

Nayani, U. (2026b). The Publisher CFO: A capital allocation framework for independent intellectual property publishers. SSRN Working Paper.

Author: Umran Nayani | Little Black Journal LLP | March 2026

© Copyright 2026 Umran Nayani All Rights Reserved. Protected with www.protectmywork.com, Reference Number: 15755020326S078