



Response to the CMA's Consultation on a Proposed Steering Conduct Requirement for Apple's Mobile Platform

A quantitative assessment of the proposed requirement's expected effect on structural lock-in, using the Black Hole Index.

Submitted by: Ivan Savich, Creator of the Black Hole Index (independent researcher)

Consultation: Proposed steering conduct requirement for Apple's mobile platform, opened 30 June 2026

Prior engagement: Third-party submission to the Call for Evidence on app store rules, 8 April 2026; receipt acknowledged by the Digital Markets Unit on 13 April 2026

Declaration: I have no commercial relationship with Apple, Google, or any party to this investigation. All data used is publicly available. This response is non-confidential and may be published in full.

Questions addressed: 6.2, 6.3, 6.4, 6.6, 6.7 and 6.8 of the consultation document

1. Introduction and standing

1.1. I welcome the opportunity to respond to the CMA's consultation on a proposed steering conduct requirement ("Steering CR") for Apple's mobile platform. This response provides a quantitative, model-based assessment of the proposed requirement's expected effect on structural lock-in, using the Black Hole Index (BHI) — an open measurement instrument that produces a single dimensionless ratio, $B = \text{Capture} / \text{Escape}$, for any platform. The methodology, all 184 current platform scores, and the full research paper are publicly available at blackholeindex.com under CC-BY 4.0.

1.2. In my submission of 8 April 2026 to the Call for Evidence, I recommended that the CMA "*prioritise substantive steering removal over procedural commitments*" (recommendation 9.1(b)) and offered to provide "*sensitivity analyses for proposed interventions*" (§10.3). The proposed Steering CR is precisely such an intervention. This response delivers the offered analysis.

2. Summary of position

2.1. **I support the proposed Steering CR.** Of the conduct requirements available to the CMA, steering rules act most directly on the parameters that constitute structural lock-in — portability and substitutability — rather than on procedural transparency.

2.2. The central quantitative finding of this response: a **substantively implemented** Steering CR is modelled to reduce the structural lock-in of Apple's App Store by approximately **12–23 per cent**, while a **procedurally frictioned** implementation — formally compliant links surrounded by dissuasive design, punitive fee structures, or narrow entitlements — is modelled to have **an effect near zero**. The design details on which the CMA is consulting are therefore not implementation minutiae; on this model they are the difference between the remedy working and not working.

2.3. These are model-implied structural effects, not causal welfare estimates. The status and limitations of the instrument are set out candidly in §7. Within the consultation’s structure, §4 of this response addresses question 6.2 (benefits of the proposed Steering CR), and §5 addresses questions 6.3, 6.4, 6.6, 6.7 and 6.8 (scope, user experience, WebViews and steering fees).

3. What has changed since April: the baseline has moved

3.1. My April submission scored Apple’s mobile ecosystem at B = 5.22 on the then-current release (112 platforms). The current Q3-2026 release (184 platforms, published 5 July 2026) scores the same entity at B = 4.60. The change follows a reassessment of one parameter: portability, from x = 3 to x = 5. The reassessment incorporates the cumulative legal and operational position of link-out steering — notably the United States requirement, in force since 30 April 2025, to permit external purchase links (Apple’s challenge to the resulting contempt ruling was accepted for review by the US Supreme Court on 30 June 2026, and litigation concerning Apple’s ability to impose a commission on link-out transactions remains ongoing), and Apple’s June 2025 revision of its EU terms. Much of that position pre-dates my April submission. The difference between 5.22 and 4.60 should therefore not be read as a structural change occurring wholly between April and July 2026: it combines movement in the external environment with a revised application of the published portability rubric to facts that were partly already present in April.

3.2. That distinction matters for monitoring design, which my April submission recommended (§9.1(c)): quarterly measurement should separately disclose changes driven by new external evidence and changes driven by reassessment under the scoring rubric — only the former constitutes an observed quarter-on-quarter change in the regulated environment, and I adopt that disclosure discipline from this cycle onward. The practical implication for the UK remedy stands either way: part of the portability improvement is already embedded in the current baseline, so the Steering CR’s incremental effect now runs primarily through substitutability of distribution channels, residual portability, and transaction-flow behaviour.

3.3. The Q3-2026 release also scores Apple’s App Store as a distinct entity — the object the proposed Steering CR directly regulates:

ENTITY	APRIL (Q1 RELEASE)	CURRENT (Q3 RELEASE)	DRIVER OF CHANGE
Apple mobile ecosystem	5.22	4.60	Portability reassessed 3 → 5 (see §3.1–3.2)
Apple App Store (distribution layer)	6.13*	5.26	Scored as a separate entity since Q4-2025; currently x = 4, s = 3

* The App Store entity carried a score in the Q1 release but was not cited in the April submission, which analysed the ecosystem level. Historical quarterly values, including the Q1 figure of 6.13, are publicly retrievable via the quarterly history API (blackholeindex.com/api/platform/apple/history) and the archived Q1-2026 data snapshot in the open repository (github.com/Joingithubstyle/blackholeindex/data).

4. Modelled effect of the proposed Steering CR

4.1. Steering restrictions bear on three of BHI’s eleven parameters, as set out in my April submission (§6.1): portability (x), because steering prohibitions prevent developers from establishing direct customer and payment relationships; substitutability (s), because forcing all transactions through the

store suppresses practical substitutability of distribution channels even where alternatives nominally exist; and closeness (c), because mandatory in-app purchase flows route additional transaction frequency through the platform.

4.2. Applying the published V3.1 model to the current Q3 baseline of the directly regulated entity, Apple's App Store (B = 5.26):

SCENARIO • APPLE APP STORE	PARAMETER CHANGES	B	CHANGE
Current state	x=4, s=3, c=8	5.26	—
Substantive Steering CR (conservative)	x+1, s+1	4.63	−12.0%
Substantive CR incl. transaction-flow effect	x+1, s+1, c−1	4.05	−23.0%
Upper-bound sensitivity	x+2, s+1, c−1	3.84	−27.0%
Frictioned / procedural implementation	no parameter change	5.26	≈ 0%

4.3. At the ecosystem level, where the April analysis was conducted, the incremental effect is smaller because part of the portability improvement has already been absorbed into the baseline (§3.1): a substantive Steering CR is modelled at −5 to −10 per cent from the current B = 4.60 (x+1 alone: 4.37; x+1 and s+1: 4.13).

4.4. For continuity: re-computing the April ecosystem scoring (x = 3) under the current model reproduces B = 5.22 exactly. The April scenario table and this response are outputs of the same published, deterministic formula.

5. The design conditions on which the modelled effect depends

5.1. The modelled gap between −12 to −23 per cent and ≈ 0 in the table above is driven primarily by implementation design, as represented in the stated parameter assumptions. My April submission warned (§6.3) that if platforms implement alternative payment options with “*significant friction, adverse user experience design, or informational asymmetries*” — as has been observed in some EU DMA compliance efforts — the effect on portability and substitutability scores would be minimal. That warning maps directly onto the design choices in this consultation:

5.2. (a) **Fee level and structure (questions 6.8(a)–(c)).** If the charge applied to steered transactions approximates the in-store commission, the economic incentive to steer disappears and the substitutability parameter does not move. The modelled −12 to −23 per cent assumes steering on terms that leave developers a material economic difference. The observable precedent is Apple's June 2025 revision of its EU terms, which replaced the formal steering prohibition with a multi-component fee structure — an outcome the consultation document itself records at §4.100(b): low adoption of steering in the EU, in part because of the resulting fee level. Against that background, three features of the proposal are well designed: the definition of Steering Fee as including “*any other such fee which, by its level, structure or operation, has an equivalent effect*”; the cost-plus principle, which — anchored to genuinely incremental costs, as §4.107(a) notes these could be very low — protects the economic differential on which the modelled effect depends; and the safeguards built into the value-based principle at paragraph 13, in particular the required adjustment for substantial and entrenched market power and the netting of developers' own contribution. The residual risk is in application rather than drafting: §4.126 recognises that value assessment is complex and open to subjectivity, and the interpreta-

tive notes leave Apple flexibility in how the principles are combined. I would encourage the CMA to hold the market-power adjustment at paragraph 13(d) strictly — quantitative measures of structural dependency, of which BHI is one open example, are instruments built for precisely that adjustment.

5.3. (b) Choice architecture (question 6.6). Interstitial warning screens, additional confirmation steps, and dissuasive language suppress the behavioural component of the effect (the c parameter) even where links are formally permitted. The developer estimates recorded in the consultation document — expected drop-off from a single neutral screen ranging from 5 per cent to over 60 per cent (§4.61) — provide relevant empirical and stakeholder evidence for this behavioural channel, though the CMA notes at §4.62 that developers are unlikely to have tested the exact proposed scenario; the heightened sensitivity of microtransaction models recorded there is what the closeness parameter would predict. A single neutral interstitial screen, as proposed, is a reasonable balance between steering effectiveness and user protection — provided it remains genuinely single and genuinely neutral. The refinements contemplated at §4.66 and question 6.6(b) — time-bounding, frequency-bounding, changing defaults, or restricting the screens to riskier transactions — would each be expected to reduce the behavioural suppression and are, on this model, worth adopting.

5.4. (c) Breadth of the entitlement (questions 6.3 and 6.7). If steering is limited to narrow link formats, excluded app categories, or restricted offer types, the substitutability change applies only to a fraction of transaction volume, reducing the aggregate effect roughly in proportion to the affected transaction share (a simplifying linear assumption). Of the scope restrictions consulted on, limiting links to the developer's own website (question 6.3(a)) is a modest constraint. Restricting links to the immediate completion of a transaction (question 6.3(b)) deserves more care: account creation, subscription management and offer-comparison journeys bear on both portability and practical substitutability even where the purchase is not completed immediately after redirection, and the free-text protections at §4.26 do not fully substitute for them. Separately, a prohibition on steering to WebViews (question 6.7) adds friction on the completion path itself — developers report WebViews provide a smoother purchase journey than external browser flows (§4.70) — and would partially offset the effect of the redirection right it accompanies.

5.5. (d) Parity of presentation. Paragraphs 6 to 8 of the proposed CR already secure this condition in the drafting: equal treatment of redirection mechanisms, a single neutral interstitial screen as the only permitted addition, and a prohibition on further frictions including at the point of price comparison. The modelled substantive scenario assumes these provisions operate as drafted and are enforced; they are, on this model, load-bearing and should be retained without dilution.

CENTRAL POINT FOR THE CONSULTATION

On this model, the Steering CR's effect on structural lock-in is not determined by whether steering is permitted, but by whether steering is *economically and behaviourally viable*. Each of the four design dimensions can materially erode the remedy's structural effect; severe friction on any single dimension could move the outcome toward the frictioned scenario in the table above.

6. Ex-post verification: how the CMA could test whether the remedy worked

6.1. A structural scenario estimate should be falsifiable after the fact, and primarily against independent outcome data rather than my own re-scoring. I commit to re-scoring Apple's mobile ecosystem and App Store quarterly, on the published rubrics, before and after the Steering CR takes effect, as a

complementary structural record; all scores and their rationales are published openly. If the observable conditions assumed in the substantive scenario are present — a material commission differential and limited behavioural friction, as evidenced by the metrics at §4.153 — and portability and substitutability nonetheless fail to improve within two to three quarters, that would count as evidence against the modelled mapping, and I will publish that result.

6.2. The compliance metrics the CMA proposes to require at §4.153 — the share of transactions via IAP versus steering, commissions paid on steered transactions disaggregated by developer size, and notably the percentage rate of abandonment at mandatory information screens — are exactly the outcome measures needed to distinguish substantive uptake from formal compliance, and I support them. Two marginal additions would strengthen the set: the effective commission differential between steered and in-store transactions (the economic incentive that drives the substitutability parameter), and ex-ante definition of what abandonment rate at the interstitial screen would trigger the refinements contemplated at §4.66.

7. Methodology status and limitations

7.1. BHI is a proposed measurement instrument at a preliminary validation stage, and this response should be weighted accordingly. Current status, all public: preliminary validation shows Spearman $\rho = 0.79$ ($p < 0.001$) against annual retention rates for a 14-platform reference cohort; inter-rater reliability is pending — all scores to date come from a single evaluator; independent multi-evaluator scoring is planned across Q3–Q4 2026, with a pre-registered V4 study targeted for Q4 2026 on OSF.

7.2. The estimates in §4 are model-implied structural effects under stated parameter assumptions — scenario analysis, not causal identification. BHI is descriptive, not causal. The instrument's sensitivity to scoring subjectivity and to the choice of aggregation function has been quantified and published, with code and a pinned data snapshot, at blackholeindex.com/robustness: rank-level results are robust to ± 1 scoring noise, while absolute B values and zone thresholds are not — which is why this response emphasises directions and relative magnitudes rather than precise point values.

7.3. Scoring of the affected parameters (x, s, c) involves anchored but ultimately judgment-based rubrics. The scenario range in §4.2, rather than a point estimate, reflects this uncertainty.

8. Recommendations

8.1. **Adopt the Steering CR, with the design dimensions in §5 treated as first-order.** Fee structure, choice architecture, entitlement breadth, and presentation parity jointly determine whether the requirement has a structural effect at all.

8.2. **Define compliance substantively, not formally.** A compliance test based on observable outcomes (§6.2) — steered-transaction share, effective commission differential — is harder to satisfy with frictioned implementations than a test based on the formal availability of links.

8.3. **Monitor structural lock-in over time.** Quarterly measurement — by BHI or any comparable instrument — can provide a transparent and reproducible longitudinal record of whether interventions reduce structural dependency or merely change its surface presentation. The disclosure approach adopted in §3 illustrates how such monitoring should distinguish observed external change from reassessment under the scoring rubric.

8.4. I remain available to provide further analysis, including scenario assessments of specific design options under consultation, or a methodology briefing for the team.

Ivan Savich

CREATOR, BLACK HOLE INDEX · INDEPENDENT RESEARCHER

ivan.savich@blackholeindex.com · blackholeindex.com · ORCID 0009-0003-9542-7022

References

Savich, I. (2026). Third-Party Submission to the CMA's Call for Evidence on Recent Developments in Relation to Apple's and Google's App Store Rules. 8 April 2026. blackholeindex.com/cma-submission

Savich, I. (2026). The Black Hole Index: A Structural Lock-In Measurement Framework for Digital Platforms. SSRN Abstract 6555158.

Black Hole Index (2026). Robustness note: scoring-noise Monte Carlo and aggregation-form ablation, Q3-2026 dataset. blackholeindex.com/robustness

Farrell, J. and Klemperer, P. (2007). 'Coordination and Lock-In: Competition with Switching Costs and Network Effects.' In *Handbook of Industrial Organization*, Vol. 3.

Katz, M. L. and Shapiro, C. (1985). 'Network Externalities, Competition, and Compatibility.' *American Economic Review*, 75(3), pp. 424–440.

Samuelson, W. and Zeckhauser, R. (1988). 'Status Quo Bias in Decision Making.' *Journal of Risk and Uncertainty*, 1(1), pp. 7–59.

Competition and Markets Authority (2026). Consultation document: proposed steering conduct requirement for Apple's mobile platform. 30 June 2026.

Digital Markets, Competition and Consumers Act 2024 (UK).