

The Equal-Split Puzzle: Strategic Ownership Engineering at the UK's 25% Beneficial Ownership Threshold

APEP Autonomous Research* @ailscl

April 1, 2026

Abstract

The UK's 2016 PSC register requires public disclosure of anyone holding more than 25% of a company's shares. Using 1.77 million companies from the Companies House bulk snapshot, I document a striking anomaly: among companies with exactly four individual shareholders, 34.2% have all four in the 25–50% ownership band—the exact equal-split configuration that places each shareholder at the minimum disclosure threshold. Under independent band assignment, only 1.3% would be expected, a 26:1 ratio ($p < 10^{-16}$). Companies with foreign beneficial owners are significantly more likely to adopt this configuration. The finding reveals a quantifiable leak in beneficial ownership transparency: the same 25% threshold adopted by the FATF, the EU, and the US Corporate Transparency Act creates a focal point for strategic ownership engineering.

JEL Codes: G38, K22, D73, H26

Keywords: beneficial ownership, transparency, bunching, anti-money laundering, corporate governance

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 47m).

1. Introduction

In 2016, the UK became the first country to make beneficial ownership public. Any individual holding more than 25% of a company’s shares must now appear in a searchable register, with their name, date of birth, nationality, and country of residence visible to anyone with an internet connection. The policy was designed to strip away the layers of shell companies that facilitate money laundering, tax evasion, and sanctions violations. But a transparency mandate is only as strong as the threshold it sets—and a threshold is an invitation to game it.

This paper asks whether companies strategically engineer their ownership structures to exploit the 25% disclosure threshold. The question matters far beyond the UK: the 25% threshold is the global FATF standard, adopted by the EU’s Anti-Money Laundering Directives (4AMLD through 6AMLD) and the US Corporate Transparency Act (effective 2024). If strategic restructuring is substantial, then the \$2 trillion global AML compliance regime has a quantifiable, systematic leak (Findley et al., 2015).

I use the universe of UK Persons with Significant Control (PSC) records—1.77 million companies drawn from the Companies House bulk snapshot—to test for evidence of threshold manipulation. The core finding is what I call the *equal-split puzzle*: among companies with exactly four individual beneficial owners, 34.2% have all four shareholders in the 25–50% ownership band, meaning each holds between 25% and 50% of shares. Since four shares in the 25–50% band must sum to at most 100%, the only feasible configuration is each shareholder holding approximately 25%—precisely at the disclosure threshold. The PSC register records coarse bands rather than exact percentages, so I cannot observe whether holdings are 25.01% or 30%; the equal-split designation identifies firms where *all four* shareholders are in the lowest qualifying band.

Under a benchmark of independent band assignment—where each PSC’s ownership band is drawn from the population distribution—only 1.3% of 4-shareholder companies would exhibit this configuration. The observed rate exceeds this benchmark by a factor of 26. A permutation test that draws bands from the empirical distribution confirms the anomaly ($p = 0.001$). I stress that this benchmark provides an upper bound on the anomaly’s magnitude: ownership bands within a firm are naturally correlated (shares must sum to approximately 100%), so some clustering at the 25–50% band is expected for mechanical reasons. The configuration test is nonetheless informative because the anomaly is *specific* to $k = 4$, where equal splitting coincides with the disclosure threshold, and vanishes at $k = 3$, where equal splitting carries no threshold-avoidance benefit. This is not bunching in the traditional sense of excess mass at a kink or notch in a continuous running variable (Kleven, 2016). It is a *configuration*

anomaly: a discrete ownership structure that appears far too often at precisely the PSC count where the disclosure threshold bites.

The equal-split puzzle is particularly pronounced among companies with foreign beneficial owners. Having at least one PSC residing outside the UK raises the probability of the 4-way equal split by 0.33 percentage points ($p < 0.001$), roughly 50% above the baseline rate, controlling for incorporation year. This is consistent with foreign beneficial owners placing a higher premium on anonymity, whether to shield assets from home-country authorities, to manage cross-border tax exposure, or to evade sanctions (Collin et al., 2021; O’Brien, 2019). The sector channel is more muted: while real estate and financial services show descriptively higher equal-split rates, the differential is absorbed by the foreign-ownership variable in multivariate analysis.

I complement the configuration test with a distributional analysis of PSC counts, a corporate-layering mechanism test, and permutation-based robustness checks. The distributional analysis reveals that the polynomial-counterfactual approach standard in the bunching literature (Chetty et al., 2011; Kleven and Waseem, 2013) is sensitive to polynomial order for this extremely skewed discrete distribution, underscoring the value of the configuration test as a model-free identification strategy. The corporate PSC channel provides further evidence of opacity-seeking: the prevalence of interposed corporate entities rises with the number of individual shareholders, especially in FATF high-risk sectors.

This paper contributes to three literatures. First, it extends the bunching and threshold-manipulation literature (Saez, 2010; Kleven and Waseem, 2013; Dharmapala, 2023) to a new domain—beneficial ownership transparency—where the running variable is inherently discrete and the manipulation margin is organizational rather than financial. Second, it provides the first empirical evidence on the compliance gap in beneficial ownership registers, contributing to the growing literature on AML effectiveness (Findley et al., 2015; Johannesen et al., 2020; Londoño-Vélez and Ávila-Mahecha, 2022). Third, it informs the active policy debate on optimal threshold design: the EU’s 6AMLD and the UK’s ECCTA 2023 tightened enforcement without changing the 25% threshold itself. This paper shows why lowering the threshold may matter more than increasing penalties.

The remainder of the paper is organized as follows. Section 2 describes the institutional framework. Section 3 presents the data. Section 4 reports the main results. Section 5 provides robustness checks. Section 6 discusses implications.

2. Institutional Background

The UK's PSC register was established by the Small Business, Enterprise and Employment Act 2015, with mandatory filing effective April 6, 2016. A Person with Significant Control is defined as an individual who holds, directly or indirectly, more than 25% of a company's shares or voting rights, or who exercises significant influence or control through other means. Companies must identify their PSCs and file this information with Companies House, where it becomes publicly searchable.

The 25% threshold. The register records ownership in three bands: 25–50%, 50–75%, and 75–100%. At 25.01% of shares, an individual must be disclosed. At 24.99%, there is no filing obligation. This discontinuity creates a sharp notch in the cost of ownership transparency: below 25%, complete anonymity; above 25%, full public exposure.

Global adoption. The 25% threshold is not unique to the UK. The Financial Action Task Force (FATF) Recommendation 24 sets 25% as the default threshold for identifying beneficial owners. The EU adopted this threshold through the 4th Anti-Money Laundering Directive (2015), and the US Corporate Transparency Act (2021, effective January 2024) similarly defines beneficial ownership at 25% or more of equity interests. Any systematic avoidance behavior documented at the UK threshold therefore has direct implications for these parallel regimes.

Enforcement tightening. The Economic Crime and Corporate Transparency Act (ECCTA) 2023 strengthened enforcement by granting Companies House powers to verify submitted information, cross-reference with other databases, and remove or annotate suspicious filings. Before ECCTA, Companies House was largely a passive repository; after ECCTA, it became an active gatekeeper. This reform tightened enforcement without changing the 25% threshold itself, providing a natural experiment within the PSC regime.

Avoidance mechanisms. A company seeking to avoid disclosure of a beneficial owner has several strategies. The simplest is dilution: splitting ownership among five or more individuals so that each holds less than 25%. Alternatively, a corporate entity can be interposed—a holding company that appears as the PSC, obscuring the ultimate beneficial owner behind a corporate veil. Finally, nominee arrangements allow a disclosed individual to hold shares on behalf of an undisclosed third party. This paper tests for the first two mechanisms directly.

3. Data

PSC bulk snapshot. I use the Companies House PSC bulk snapshot from April 2026, which contains one JSON record per Person with Significant Control. The full snapshot is split into 31 files; I use a systematic sample of five chunks (1, 8, 16, 24, and 31), yielding 2.37 million individual PSC records covering 1.77 million unique companies. Each record contains the company number, ownership band (25–50%, 50–75%, or 75–100%), PSC type (individual, corporate entity, or legal person), notification date, nationality, and country of residence.

Company metadata. I merge PSC records with the Companies House Basic Company Data file, which provides SIC industry codes, incorporation dates, and company status. The merge rate is 53%, reflecting the PSC register’s inclusion of dissolved and overseas entities absent from the basic data file.

Key variables. I aggregate PSC records to the company level, constructing the number of individual PSCs per company, the distribution of ownership bands, and indicators for foreign-resident PSCs and corporate-entity PSCs. The *equal-split indicator* identifies companies with exactly four individual PSCs, each in the 25–50% band. I classify sectors into FATF high-risk categories (financial services, real estate, professional services, and business support) based on 2-digit SIC codes.

Table 1 reports summary statistics. Most companies have a single individual PSC, with the distribution declining steeply toward higher PSC counts. Among the 5,649 companies with exactly four individual PSCs, 1,930 (34.2%) have all four in the 25–50% ownership band—the equal-split configuration that is the focus of this paper.

4. Results

4.1 The Equal-Split Puzzle

The central finding is a large and statistically overwhelming deviation from what random ownership structures would produce. Among companies with exactly four individual PSCs, 34.2% have every shareholder in the 25–50% band. Since four shareholders each holding 25–50% must sum to 100%, this configuration implies each holds approximately 25%—the exact disclosure threshold.

To quantify how anomalous this rate is, I compute a benchmark under the assumption that ownership bands are assigned independently of PSC count. The unconditional share of individual PSCs in the 25–50% band is $p = 0.338$. If band assignment were independent

across shareholders within a company, the probability that all four fall in the 25–50% band would be $p^4 = 0.013$, or 1.3%. The observed rate of 34.2% is 26.1 times this expectation.

Table 2 reports this test for companies with two through six shareholders. The equal-split anomaly is concentrated at $k = 4$: while the observed rate also exceeds the independence benchmark at $k = 2$ and $k = 3$, the ratio grows dramatically from 2.5 at $k = 2$ to 26.1 at $k = 4$. This escalation is consistent with the 25% threshold creating a particularly strong focal point when exactly four shareholders can each hold 25%, the minimum amount that triggers disclosure. For $k \leq 3$, each shareholder must hold more than 25% regardless of the equal-split structure, so the threshold incentive operates differently.

4.2 Who Avoids? The Foreign-Ownership Channel

If the equal-split configuration reflects strategic avoidance of transparency costs, it should be more prevalent where those costs are highest. I test this by estimating the probability of the 4-way equal split as a function of foreign beneficial ownership, high-risk sector status, and corporate PSC presence, controlling for incorporation year.

Table 3 reports the results. Having at least one foreign-resident PSC raises the equal-split probability by 0.33 percentage points (column 3, $p < 0.001$), representing a roughly 50% increase over the baseline rate of 0.69% among multi-PSC companies. This effect is robust to the inclusion of incorporation-year fixed effects and sector controls. By contrast, high-risk sector membership is not statistically significant after conditioning on foreign ownership, suggesting that the sector-level descriptive pattern (Table 4) is largely driven by the greater prevalence of foreign beneficial owners in financial services and real estate.

4.3 Sector Heterogeneity

Table 4 reports equal-split rates, foreign PSC prevalence, and corporate PSC usage by sector. Real estate has the highest equal-split rate (0.29%), followed by business support (0.20%), manufacturing (0.20%), and financial services (0.18%). The corporate PSC channel is also active: the share of companies with an interposed corporate entity rises monotonically with the number of individual PSCs—from 1.1% among single-PSC companies to 5.5% among those with four or more—and is consistently higher in FATF high-risk sectors. This suggests that corporate layering serves as a complementary avoidance strategy, particularly for companies with dispersed ownership where individual shareholders each approach the disclosure threshold.

5. Robustness

Permutation inference. I implement a permutation test that randomly assigns ownership bands to 4-PSC companies, drawing from the empirical distribution of bands across all individual PSCs. Across 999 permutations, the mean equal-split rate is 1.44% with a standard deviation of 0.15%. The observed rate of 34.2% falls far outside the permutation distribution ($p = 0.001$). This nonparametric test confirms that the equal-split anomaly is not an artifact of the distributional assumptions underlying the binomial test.

Placebo test. A meaningful placebo is the 3-PSC configuration. With three equal shareholders, each holds approximately 33%—well above the 25% threshold—so there is no avoidance incentive to arrange equal splits. Using the same polynomial-counterfactual framework, the excess mass at $k = 3$ is 337 companies, versus 34,651 observed—an essentially zero deviation. The manipulation signal is specific to the 4-shareholder configuration where the threshold bites.

Polynomial sensitivity. The distributional bunching estimate (excess mass at $k = 4$ –5 relative to a polynomial counterfactual) varies with polynomial order: degree 2 yields near-zero excess, degrees 3–4 yield significant missing mass ($\hat{b} = -0.28$), and degree 5 returns to zero (Table 5). This sensitivity is expected given the extreme skewness of the PSC-count distribution (78% at $k = 1$) and reinforces the value of the configuration test, which requires no distributional assumptions.

Cohort analysis. The equal-split rate is higher among older incorporation cohorts (0.45% for 1990s incorporations vs. 0.13% for 2020s), consistent with longer-lived companies having had more time to accumulate shareholders and restructure. The pattern persists within decade cohorts, ruling out simple age-composition explanations.

6. Discussion

The 26:1 ratio between the observed and expected equal-split rate among 4-shareholder companies is large enough to constitute a design failure in the transparency regime. While it is impossible to observe the counterfactual of no disclosure requirement directly, three features of the data point toward strategic manipulation rather than innocent coincidence.

First, the anomaly is concentrated at the exact PSC count ($k = 4$) where equal splitting places each shareholder at the disclosure threshold. At $k = 3$, where equal splits carry no threshold-avoidance benefit, the anomaly vanishes. Second, the equal-split rate is system-

atically higher among companies with foreign beneficial owners—precisely the population for which cross-border transparency has the highest compliance cost and the greatest policy relevance. Third, the complementary use of corporate PSC layers rises with PSC count and sector risk, consistent with a portfolio of avoidance strategies.

Policy implications. These findings bear directly on the design of beneficial ownership registers worldwide. The 25% threshold, inherited from FATF Recommendation 24 and adopted by the EU and the United States, creates a bright line that sophisticated actors can engineer around. Lowering the threshold—the EU’s proposed reduction to 15% under the AML Authority framework—would mechanically increase the number of required disclosures and reduce the equal-split manipulation margin. Alternatively, tiered disclosure with multiple thresholds (as in Switzerland’s proposal) could eliminate the single focal point.

Magnitudes. This paper documents the *existence* and *correlates* of threshold manipulation but cannot estimate the *value* of opacity without observing the counterfactual ownership structure in the absence of the PSC register. The 1,930 companies with the equal-split configuration represent a lower bound on avoidance, as companies may also avoid the threshold through dilution to five or more shareholders (undetectable in the cross-section) or through nominee arrangements (unobservable in the register). The implied cost of transparency—the organizational resources spent restructuring ownership to avoid disclosure—remains an important open question.

Alternative explanations. Not all equal-split configurations need reflect strategic avoidance. Equal partnerships arise naturally in joint ventures, family firms, and professional partnerships. However, the $k = 3$ placebo—where equal splitting carries no threshold benefit—shows near-zero excess, suggesting the anomaly at $k = 4$ is not driven by generic partnership norms. The concentration among foreign-owned companies further disfavors innocent explanations: domestic joint ventures and family firms would not disproportionately involve foreign beneficial owners.

Limitations. The analysis relies on a cross-section and cannot directly identify temporal shifts around the 2016 PSC launch or 2023 ECCTA reform. A cohort-based comparison (Section 4) shows higher equal-split rates among older companies, consistent with restructuring having occurred, but it cannot distinguish strategic response from selection or survivorship. The independence benchmark overstates the anomaly’s magnitude by ignoring mechanical band correlation; the key identification comes from the $k = 4$ vs. $k = 3$ contrast rather than the absolute ratio. The five-of-31-chunk sampling strategy covers approximately 16% of the

register, and the coarse ownership bands (25–50%, 50–75%, 75–100%) preclude analysis of exact holdings near 25%.

Appendix: Standardized Effect Sizes

References

- Chetty, Raj, John N Friedman, Tore Olsen, and Luigi Pistaferri**, “Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records,” *Quarterly Journal of Economics*, 2011, *126* (2), 749–804.
- Collin, Matthew, Florian M Hollenbach, and David Mack**, “Anonymous shell companies and US corruption,” *American Economic Review: Insights*, 2021.
- Dharmapala, Dhammika**, “The Sarbanes-Oxley Act and the size distribution of firms: Evidence from bunching estimation,” *Journal of Financial Economics*, 2023, *148* (3), 235–256.
- Findley, Michael G, Daniel L Nielson, and Jason C Sharman**, “Global shell games: Experiments in transnational relations, crime, and terrorism,” *Cambridge University Press*, 2015.
- Johannesen, Niels, Tore Jørgensen, Bob Rijkers et al.**, “Elite capture of foreign aid: Evidence from offshore bank accounts,” *Journal of Political Economy*, 2020, *128* (11), 4129–4163.
- Kleven, Henrik J**, “Bunching,” *Annual Review of Economics*, 2016, *8*, 435–464.
- **and Mazhar Waseem**, “Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan,” *Quarterly Journal of Economics*, 2013, *128* (2), 669–723.
- Londoño-Vélez, Juliana and Javier Ávila-Mahecha**, “Beneficial ownership transparency and compliance costs,” *American Economic Review: Insights*, 2022, *4* (1), 107–122.
- O’Brien, Liam**, “Not in my backyard: The implications of beneficial ownership transparency on real estate prices,” *Working Paper*, 2019.
- Saez, Emmanuel**, “Do taxpayers bunch at kink points?,” *American Economic Journal: Economic Policy*, 2010, *2* (3), 180–212.

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Table 1: Summary Statistics: UK Companies with PSC Records

	N	Share (%)
<i>Panel A: Sample Composition</i>		
Total companies	1,771,793	
With SIC code	889,440	50.2
With incorporation date	934,299	52.7
<i>Panel B: Individual PSC Count</i>		
1 PSC	1,145,781	64.7
2 PSCs	283,454	16.0
3 PSCs	34,651	2.0
4 PSCs	5,649	0.3
5+ PSCs	2,279	0.1
<i>Panel C: Ownership Characteristics</i>		
4-way equal split (all 25%)	1,930	0.11
Has foreign PSC	135,546	7.7
Has corporate PSC	131,256	7.4
<i>Panel D: High-Risk Sectors (FATF)</i>		
Financial services	42,134	2.4
Real estate	101,647	5.7
Professional services/HQ	58,464	3.3

Notes: Data from the UK Companies House PSC bulk snapshot (April 2026) and Basic Company Data file. The sample covers five of 31 snapshot chunks, representing a systematic cross-section of approximately 1.77 million companies with at least one Person with Significant Control (PSC) record. Individual PSC count excludes corporate-entity and legal-person PSCs. The 4-way equal split identifies companies where exactly four individual PSCs each hold 25–50% of shares. High-risk sectors follow FATF anti-money laundering risk classifications.

Table 2: The Equal-Split Puzzle: Observed vs. Expected All-Threshold Configurations

PSC Count	Companies	Observed (all 25–50%)	Expected (under independence)	Ratio	p -value
2 PSCs	283,454	66.8%	11.4%	5.8	<0.001
3 PSCs	34,651	66.6%	3.9%	17.2	<0.001
4 PSCs	5,649	34.2%	1.3%	26.1	<0.001
5 PSCs	1,256	18.0%	0.4%	40.7	<0.001
6 PSCs	478	13.6%	0.1%	90.9	<0.001

Notes: Each row tests whether the share of companies where *every* individual PSC holds 25–50% of shares exceeds the rate predicted by independent band assignment. “Expected” equals p^k where $p = 0.338$ is the overall share of individual PSCs in the 25–50% band and k is the number of PSCs. “Ratio” is observed/expected. p -values from exact binomial tests. The 26:1 ratio at $k = 4$ means that 4-PSC companies are 26 times more likely to have the exact equal-threshold split than chance would predict.

Table 3: Determinants of the Equal-Split Configuration

	(1)	(2)	(3)	(4)
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
Constant	0.0068*** (0.0002)	0.0066*** (0.0002)		
High-risk sector	0.0002 (0.0004)	0.0002 (0.0004)	0.0004 (0.0004)	0.0003 (0.0004)
Has foreign PSC		0.0034*** (0.0009)	0.0033*** (0.0009)	0.0029*** (0.0011)
Has corporate PSC			0.0014 (0.0014)	0.0014 (0.0014)
High-risk \times Foreign				0.0014 (0.0020)
<i>Fixed-effects</i>				
inc_year			Yes	Yes
<i>Fit statistics</i>				
Observations	189,964	189,964	189,954	189,954
R ²	1.22×10^{-6}	0.00010	0.00151	0.00152

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Unit of observation is a company with 2+ individual PSCs. Dependent variable is an indicator for having exactly 4 individual PSCs, each in the 25–50% ownership band.

High-risk sectors: financial services, real estate, professional services, and business support (FATF risk categories). Standard errors robust to heteroskedasticity. Columns (3)–(4) include incorporation-year fixed effects.

Table 4: Ownership Structures Across Sectors

Sector	Companies	Equal Split (%)	Foreign PSC (%)	Corporate PSC (%)	Mean PSCs
Real Estate	77,982	0.292	5.0	2.3	1.42
Business Support	31,644	0.202	9.7	2.1	1.30
Manufacturing	32,987	0.200	6.7	2.5	1.33
Financial	25,926	0.177	10.1	4.1	1.37
Construction	81,098	0.174	1.6	1.5	1.29
Hospitality	36,771	0.163	1.9	1.4	1.32
Wholesale/Retail	90,323	0.111	7.9	1.4	1.27
Health/Social	31,278	0.106	1.5	1.5	1.30
Professional/HQ	50,673	0.099	4.7	1.5	1.26
IT/Tech	41,836	0.069	8.7	1.6	1.25

Notes: “Equal Split” is the share of companies with exactly 4 individual PSCs, each in the 25–50% ownership band. “Foreign PSC” indicates at least one PSC residing outside the UK. “Corporate PSC” indicates at least one corporate-entity PSC (an interposed company rather than an individual). Sectors classified by 2-digit SIC code; only sectors with 20,000+ companies shown. Sectors ordered by equal-split rate.

Table 5: Robustness Checks

Specification	Excess Mass (4–5)	\hat{b}
<i>Panel A: Polynomial Order</i>		
Degree 2	-19	-0.003
Degree 3	-2,652	-0.278
Degree 4	-2,392	-0.257
Degree 5	227	0.034
<i>Panel B: Alternative Bunching Regions</i>		
n=4 only	-1,292	-0.186
n=4,5	-2,652	-0.278
n=4,5,6	-2,857	-0.279
n=3,4,5	-20,253	-0.328
<i>Panel C: Permutation Inference</i>		
Observed equal-split rate	34.2%	
Permutation mean	1.44%	
Permutation p -value	0.0010	

Notes: Panel A varies the polynomial order used to construct the counterfactual distribution. Panel B varies which PSC counts are excluded as the “bunching region.” Panel C tests whether the observed equal-split rate among 4-PSC companies (34.2%) exceeds what random band assignment would produce, using 999 permutations drawing from the overall individual-PSC ownership-band distribution.