Do Community Police Matter? Evidence from England's PCSO Austerity Cuts

APEP Autonomous Research* @ai1scl

March 5, 2026

Abstract

Between 2010 and 2024, England's police forces cut their Police Community Support Officers (PCSOs)—uniformed civilian staff who patrol neighbourhoods—by approximately 60%, with enormous cross-force variation ranging from total elimination to modest expansion. I exploit this variation in a two-way fixed effects framework across 41 police force areas over 2008–2024 to estimate the causal effect of community policing on recorded crime. The point estimate is a precise zero: each additional PCSO per 100,000 population is associated with a -0.02% change in crime (SE $= 0.22\%$, $p = 0.92$). This null is robust to wild cluster bootstrap ($p = 0.93$), randomization inference ($p = 0.675$), jackknife sensitivity, and crime-type decomposition. The design can rule out crime effects larger than 9.6% from the average PCSO decline—a meaningful bound on the returns to community-oriented civilian policing.

JEL Codes: K42, H76, H75

Keywords: police, community policing, PCSOs, crime, austerity, England

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

England cut more than half its community police officers in a decade, and nobody noticed. Between 2010 and 2024, Police Community Support Officers—uniformed civilian staff who walk beats, engage with residents, and gather neighbourhood intelligence—fell from 16,000 to fewer than 7,000 nationally. Some forces eliminated PCSOs entirely; others preserved or even expanded their numbers. Yet throughout this dramatic variation, there has been no systematic effort to estimate whether these cuts affected crime.

This matters because community policing is central to modern police strategy worldwide. The idea that uniformed officers on foot patrol can deter crime, build public trust, and generate intelligence has shaped police staffing decisions from New York to Tokyo (Goldstein, 1990; Tyler, 2004). In England, PCSOs were introduced in 2002 precisely to fill this community-facing role, freeing sworn officers for response and investigation (Crawford et al., 2005). Yet PCSOs lack arrest powers. They cannot stop and search. Their toolbox is presence, conversation, and community engagement. Whether this toolbox reduces crime is an empirical question that prior work has not answered at scale.

This paper provides the first national-scale causal estimate. I exploit cross-force variation in PCSO cuts driven by England's post-2010 austerity programme, which reduced central government grants to police forces by approximately 20% in real terms (House of Commons Library, 2023). Because forces differed in their reliance on central grants, their fiscal capacity to protect PCSOs varied substantially. I estimate two-way fixed effects (TWFE) models of recorded crime rates on PCSO staffing per 100,000 population, controlling for sworn officer levels and absorbing force and year fixed effects. The identification leverages continuous dose-response variation across 41 forces over 17 years (2008–2024), yielding 697 force-year observations.

The main finding is a precise null. The preferred specification—log crime rate on PCSOs per 100,000, controlling for sworn officers—yields a coefficient of -0.0002 with a standard error of 0.0022. This estimate is economically negligible and statistically indistinguishable from zero ($p = 0.92$). A log-log specification yields a PCSO-crime elasticity of 0.027 (SE = 0.036), implying that a 10% reduction in PCSOs is associated with a 0.3% increase in crime—well within confidence intervals that span zero.

I subject this null to extensive robustness testing. Wild cluster bootstrap inference with Webb weights, designed for settings with few clusters (Cameron et al., 2008; Roodman et al., 2019), yields a p-value of 0.93. Randomization inference, which permutes PCSO levels across forces within each year (Fisher, 1935; Young, 2019), produces a p-value of 0.675. A leave-one-out jackknife shows that no single force drives the result; the coefficient ranges from

-0.002 to 0.001 across all 41 specifications. The null survives dropping the Metropolitan Police, dropping all London forces, and a first-differenced specification. An event study using baseline PCSO exposure interacted with year dummies confirms flat pre-trends and no systematic post-treatment divergence.

The decomposition by crime type offers a mechanism test. If PCSO presence deters street-level crime through visible patrol, one would expect effects concentrated in burglary, criminal damage, and public order offences—the crime types most plausibly affected by neighbourhood presence. Instead, I find null effects across all ten offence groups examined, with no coefficient individually significant and all confidence intervals comfortably spanning zero. The largest point estimates are for miscellaneous crimes against society (-0.0045) and drug offences (-0.0035), but neither is robust to multiple-testing adjustments.

Finally, I quantify what the design can and cannot detect. The minimum detectable effect at 80% power is 9.6% of baseline crime from the average 15.3 per 100,000 PCSO decline. Given that England experienced a roughly 60% decline in PCSOs nationally, this bound implies the design can rule out crime effects exceeding approximately 5% of baseline rates per 10 percentage points of PCSO reduction. While this cannot rule out small effects, it substantially constrains the returns to community-oriented civilian policing.

This paper contributes to three literatures. First, it adds to the economics of policing, which has established that *sworn* officers reduce crime (Levitt, 1997; Chalfin and McCrary, 2018; Mello, 2019; Draca et al., 2011; Klick and Tabarrok, 2005). The consistent finding from this literature is that additional police presence deters criminal activity, with elasticities ranging from -0.3 to -1.0 (Chalfin et al., 2022). Yet nearly all this evidence concerns officers with full enforcement powers. Whether civilian staff—who lack arrest authority—produce similar deterrence is unknown. My null result suggests that the crime-reducing effect of police may derive primarily from enforcement capacity rather than mere uniformed presence.

Second, this paper speaks to the community policing literature in criminology, which has emphasized the role of procedural justice, trust-building, and intelligence-gathering in crime reduction (Tyler, 2004; Bradford, 2014; Sampson et al., 1997). The theory predicts that community engagement improves police legitimacy and cooperation, generating intelligence that prevents crime. My results do not reject this channel but suggest that any such effects are too small to detect at the force-area level, even with substantial workforce variation.

Third, this paper contributes to the literature on austerity and public services in England. d'Este and Harvey (2024) estimate that overall police budget cuts increased crime, finding effects concentrated in less visible crime categories. My results complement theirs by isolating the PCSO component: while total police resources may matter, the civilian community-policing tier appears to contribute little to measurable crime outcomes. This distinction has direct implications for optimal police workforce composition.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on PCSOs and the austerity programme. Section 3 describes the data. Section 4 presents the empirical strategy. Section 5 reports results. Section 6 discusses mechanisms and implications. Section 7 concludes.

2. Institutional Background

2.1 Police Community Support Officers

Police Community Support Officers were introduced by the Police Reform Act 2002 as a new tier of the police workforce. Unlike sworn police constables, PCSOs do not hold the Office of Constable and lack full police powers. They cannot arrest suspects, conduct stop-and-search, or carry out forced entries. Their designated powers are limited to issuing fixed penalty notices, confiscating alcohol from minors, requiring names and addresses, and using reasonable force for detention pending a constable's arrival (Crawford et al., 2005).

The rationale for PCSOs was neighbourhood policing. They were designed to provide a visible, approachable presence in communities—walking beats, attending community meetings, building relationships with residents, and gathering low-level intelligence about antisocial behaviour and local crime patterns (HMIC, 2016). By freeing sworn officers from routine patrol and community engagement duties, PCSOs were intended to allow constables to focus on response, investigation, and enforcement.

Between 2002 and 2010, PCSO numbers expanded rapidly, largely funded through central government grants. By March 2010, there were approximately 16,000 PCSOs across England and Wales, representing roughly 10% of the total police workforce. The geographic distribution was uneven from the start: some forces, particularly those in rural areas, invested heavily in PCSOs as a cost-effective way to maintain community presence, while others relied more on sworn officers.

2.2 Austerity and the PCSO Decline

The Conservative-Liberal Democrat coalition government that took office in May 2010 implemented substantial reductions in public spending. Police forces in England and Wales faced real-terms cuts of approximately 20% to their central government grant over the 2010–2015 spending review period (House of Commons Library, 2023). While Police and Crime Commissioners (PCCs, elected from 2012) could partially offset grant reductions through increases in the council tax precept, this flexibility was capped by referendum thresholds.

PCSOs bore a disproportionate share of workforce reductions. Because they were typically funded from central grants rather than core budgets, and because they lacked the statutory protections enjoyed by sworn officers, chief constables found PCSOs the most administratively straightforward posts to cut. Between 2010 and 2024, national PCSO numbers fell from approximately 16,000 to under 7,000—a decline exceeding 50%.

Crucially, the magnitude of PCSO cuts varied enormously across force areas. This variation reflects differences in initial PCSO numbers, grant dependence, PCC priorities, and local fiscal capacity. Figure 2 documents this heterogeneity. Norfolk eliminated its entire PCSO workforce (from 43 per 100,000 to zero). The City of London and Cambridgeshire cut PCSOs by over 90%. At the other extreme, Dyfed-Powys *expanded* its PCSO workforce by 48%, and several forces maintained PCSOs at over 80% of their 2010 levels.

This cross-force variation in PCSO changes, conditional on overall workforce trends and fixed force characteristics, provides the identifying variation for this study. The key question is whether forces that cut PCSOs more aggressively experienced different crime trajectories than those that preserved community policing capacity.

2.3 The PCSO Role in Practice

Understanding what PCSOs actually do is essential for interpreting the null result. Their daily duties differ fundamentally from those of sworn officers. A typical PCSO shift involves foot patrol through a designated neighbourhood, checking on vulnerable residents, visiting shops and businesses, attending school talks, and filing intelligence reports on suspicious activity. They are the most visible police presence in many communities, spending the majority of their time outdoors and on foot, compared to response officers who spend much of their time in vehicles or at police stations (HMICFRS, 2019).

PCSOs serve as a bridge between the public and the police. They attend Neighbourhood Watch meetings, liaise with local councillors, and coordinate with housing associations and social services. In areas with high ethnic diversity, PCSOs have been recruited specifically for language skills and cultural knowledge. Several forces used PCSOs as the primary point of contact for community intelligence, feeding information into tasking and coordination meetings that directed sworn officer deployments.

The theory of change for PCSOs rests on three mechanisms. First, *visible deterrence*: the presence of a uniformed figure on the street raises the perceived risk of detection for potential offenders, even if the PCSO cannot make an arrest (Nagin, 2013). Second, *community intelligence*: regular engagement with residents generates soft intelligence about emerging crime patterns, drug dealing locations, and individuals of concern, which feeds into directed enforcement by sworn officers. Third, *procedural justice*: positive interactions with PCSOs build public trust in the police, increasing willingness to report crime and cooperate with investigations (Tyler, 2004; Bradford, 2014).

Each mechanism implies different crime effects. Visible deterrence should reduce opportunistic street crime (theft, criminal damage, public order). Community intelligence should reduce crime types that depend on local knowledge to detect and disrupt (drug dealing, antisocial behaviour). Procedural justice effects may increase *reported* crime (if more trusting communities report more) while reducing *actual* crime through improved cooperation. The crime type decomposition in Section 5 tests these channels.

2.4 The Broader Austerity Context

PCSO cuts did not occur in isolation. The austerity programme affected the entire criminal justice system. Probation services were restructured under the Transforming Rehabilitation programme. Courts faced staffing reductions. Youth offending teams lost funding. Local authority services—including street lighting, youth services, and drug treatment programmes—were cut simultaneously.

This broader context has two implications for identification. First, if forces that cut PCSOs also experienced larger cuts to other crime-relevant services, the PCSO coefficient could capture the combined effect of multiple austerity cuts rather than the PCSO-specific effect. Controlling for sworn officer levels partially addresses this concern, and the force fixed effects absorb time-invariant cross-force differences in fiscal capacity. Second, if the austerity programme had offsetting effects—for example, reduced probation services increasing reoffending while reduced youth services reduced the size of potential offender cohorts—the net effect on crime could be small regardless of policing changes.

The post-2019 period introduces additional complexity. The Conservative government's "Police Uplift Programme" aimed to recruit 20,000 additional sworn officers by March 2023. This programme restored total officer numbers to near pre-austerity levels in many forces. However, PCSO numbers continued to decline during the uplift period, as forces substituted sworn officers for PCSOs. This substitution—replacement of community civilians with enforcement-capable officers—may have offset any crime effect of PCSO losses in recent years.

3. Data

3.1 Police Workforce

I obtain police workforce data from the Home Office's open data tables (Home Office, 2025b), which provide full-time equivalent (FTE) counts by force, rank, worker type (officers, PCSOs, staff), and demographic characteristics for every March 31 from 2007 to 2025. The analysis panel uses data from 2008 to 2024 (see Section 3.4 for exclusions). The data distinguish PCSOs from sworn police officers and other police staff, enabling clean measurement of the community policing workforce.

I standardize force names across the 19-year panel, accounting for merges and renaming (e.g., Hampshire and Isle of Wight Constabulary). The raw workforce dataset contains 33,528 observations across 44 force areas.

3.2 Recorded Crime

Recorded crime data come from the Home Office's open data tables (Home Office, 2025a), which report counts of recorded criminal offences by police force area, offence group, and financial year. I use three separate releases covering 2002/03–2006/07, 2007/08–2011/12, and 2012/13 onwards. All three files report crime at the force-quarter-offence level; I aggregate to annual totals by force and offence group.

The data cover ten major offence groups: violence against the person, sexual offences, robbery, theft offences, criminal damage and arson, drug offences, possession of weapons offences, public order offences, miscellaneous crimes against society, and fraud offences. Fraud data are available only for 2008–2013 (most forces stopped recording fraud locally after reporting was centralized under Action Fraud), so the fraud panel is shorter than the other nine categories. I construct total crime as the sum across all categories and compute crime rates per 100,000 population.

3.3 Population

National population estimates come from the Office for National Statistics via the NOMIS API (Office for National Statistics, 2025). I allocate national population to force areas using each force's share of total sworn officer FTE in 2010, held constant over the panel (see Appendix A.3 for the formula). This avoids a mechanical relationship between time-varying officer counts and the population denominator. The resulting force-area population estimates vary only with national population growth, not with local staffing decisions.

3.4 Panel Construction

I merge workforce, crime, and population data by standardized force name and year, restricting to complete financial years. Throughout this paper, "year" refers to the calendar year in which a financial year begins: financial year 2007/08 is "2008," financial year 2023/24 is "2024." The 2007 observation is excluded because crime data begin in financial year 2007/08. The 2025 observation is excluded because crime data cover only the first quarter, making annualized crime rates unreliable. The resulting analysis panel contains 697 force-year observations across 41 matched police force areas from 2008 to 2024. Two Welsh forces (North Wales and South Wales) are excluded due to naming inconsistencies in the crime data. The panel is strongly balanced: every included force appears in all 17 years.

Key variables are:

- **PCSO per 100,000:** PCSO FTE divided by force-area population, multiplied by 100,000.
- **Officer per 100,000:** Sworn police officer FTE divided by force-area population, multiplied by 100,000.
- **Crime rate:** Total recorded offences divided by population, multiplied by 100,000.
- **Log crime rate:** Natural logarithm of the crime rate.
- **PCSO baseline:** PCSO per 100,000 in the force's 2010 observation, used for the event study.

3.5 Summary Statistics

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max
PCSOs per 100k	22.09	8.28	0.00	43.65
Officers per 100k	235.29	24.84	171.61	297.53
Total crime rate per 100k	8284.21	2119.92	1098.83	13686.36
Log crime rate	8.98	0.33	7.00	9.52
Population (millions)	1.35	2.01	0.32	13.72

Table 1 presents summary statistics for the analysis panel. The panel spans 2008–2024, covering the full pre-austerity baseline (2008–2010) and 14 years of post-austerity variation (2011–2024). The mean PCSO level is 22.1 per 100,000 with a standard deviation of 8.3, reflecting substantial cross-force and over-time variation. Sworn officers average 236 per 100,000—roughly eleven times the PCSO level—confirming that PCSOs are a small fraction of total police resources. The mean crime rate is 8,284 per 100,000, with considerable dispersion (SD = 2,120).

Figure 1 shows national trends. Panel A documents the PCSO decline: the population-weighted mean fell from 30.7 per 100,000 in 2010 to 12.2 in 2024, a 60% reduction over 14 years. Panel B shows that total crime followed a different trajectory: declining from 2008 to 2013, rising through 2019, dipping during COVID-19, and rising again through 2024. The absence of a clear co-movement between PCSO levels and crime rates foreshadows the null result.

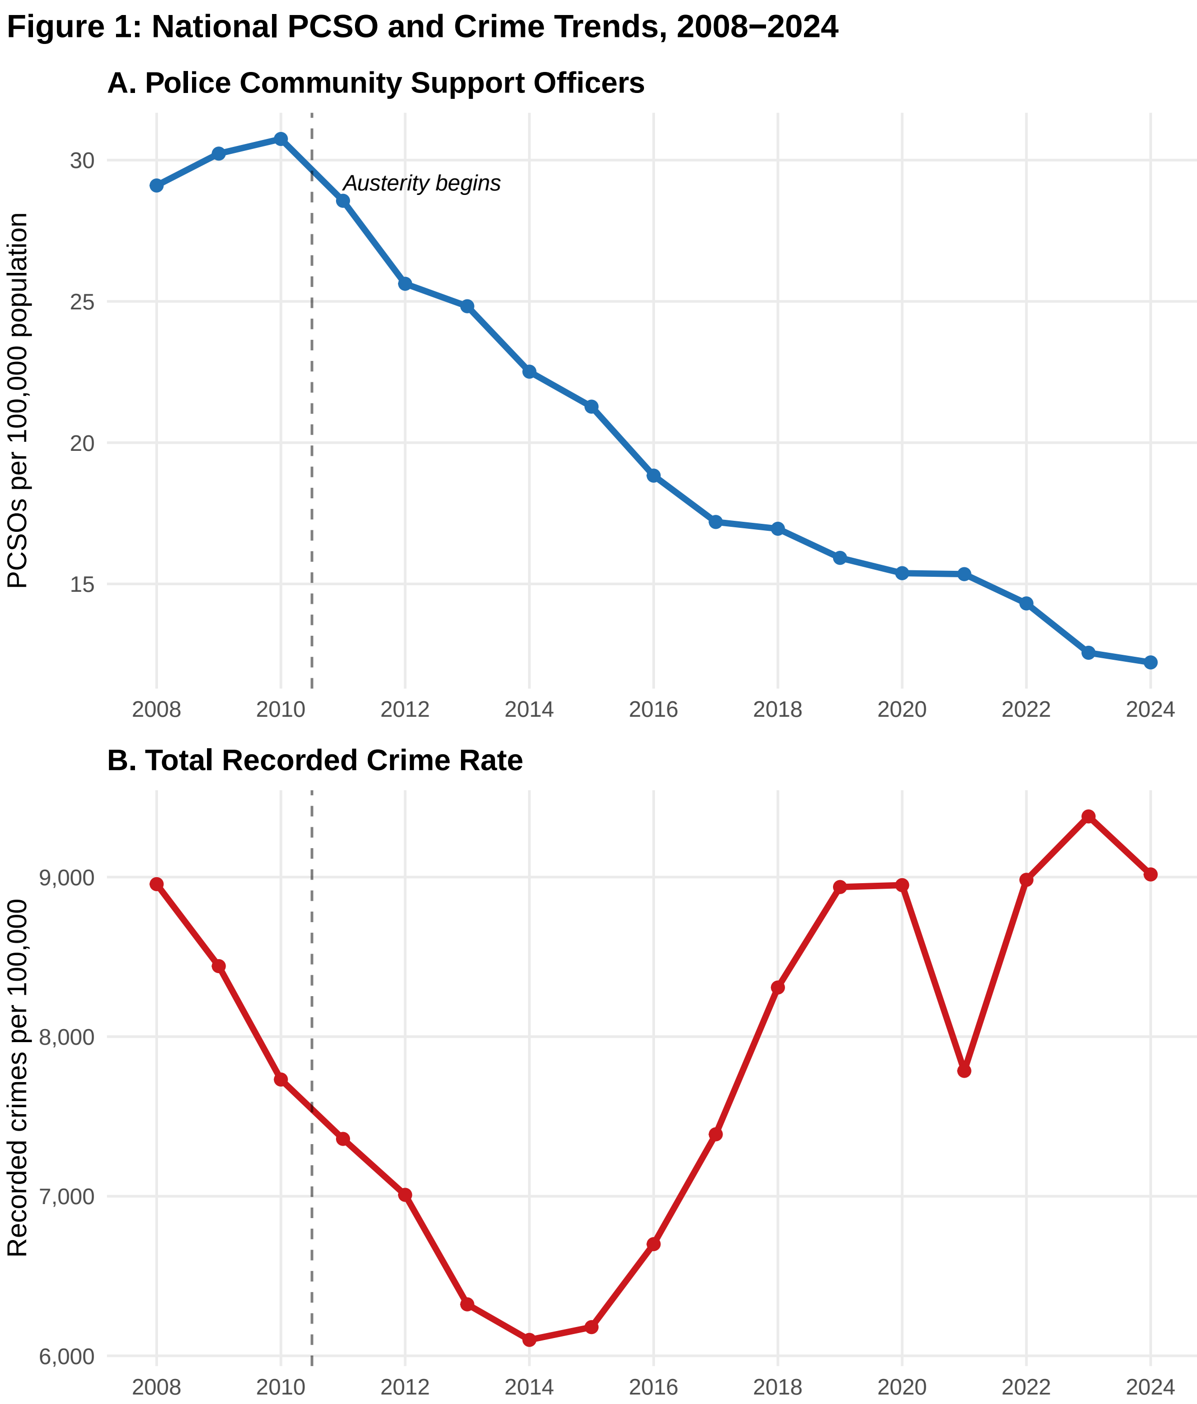

Figure 1: National PCSO and Crime Trends, 2008–2024

Notes: Population-weighted means across 41 police force areas. Dashed line marks the beginning of austerity (2010). Crime rate is recorded offences per 100,000 population.

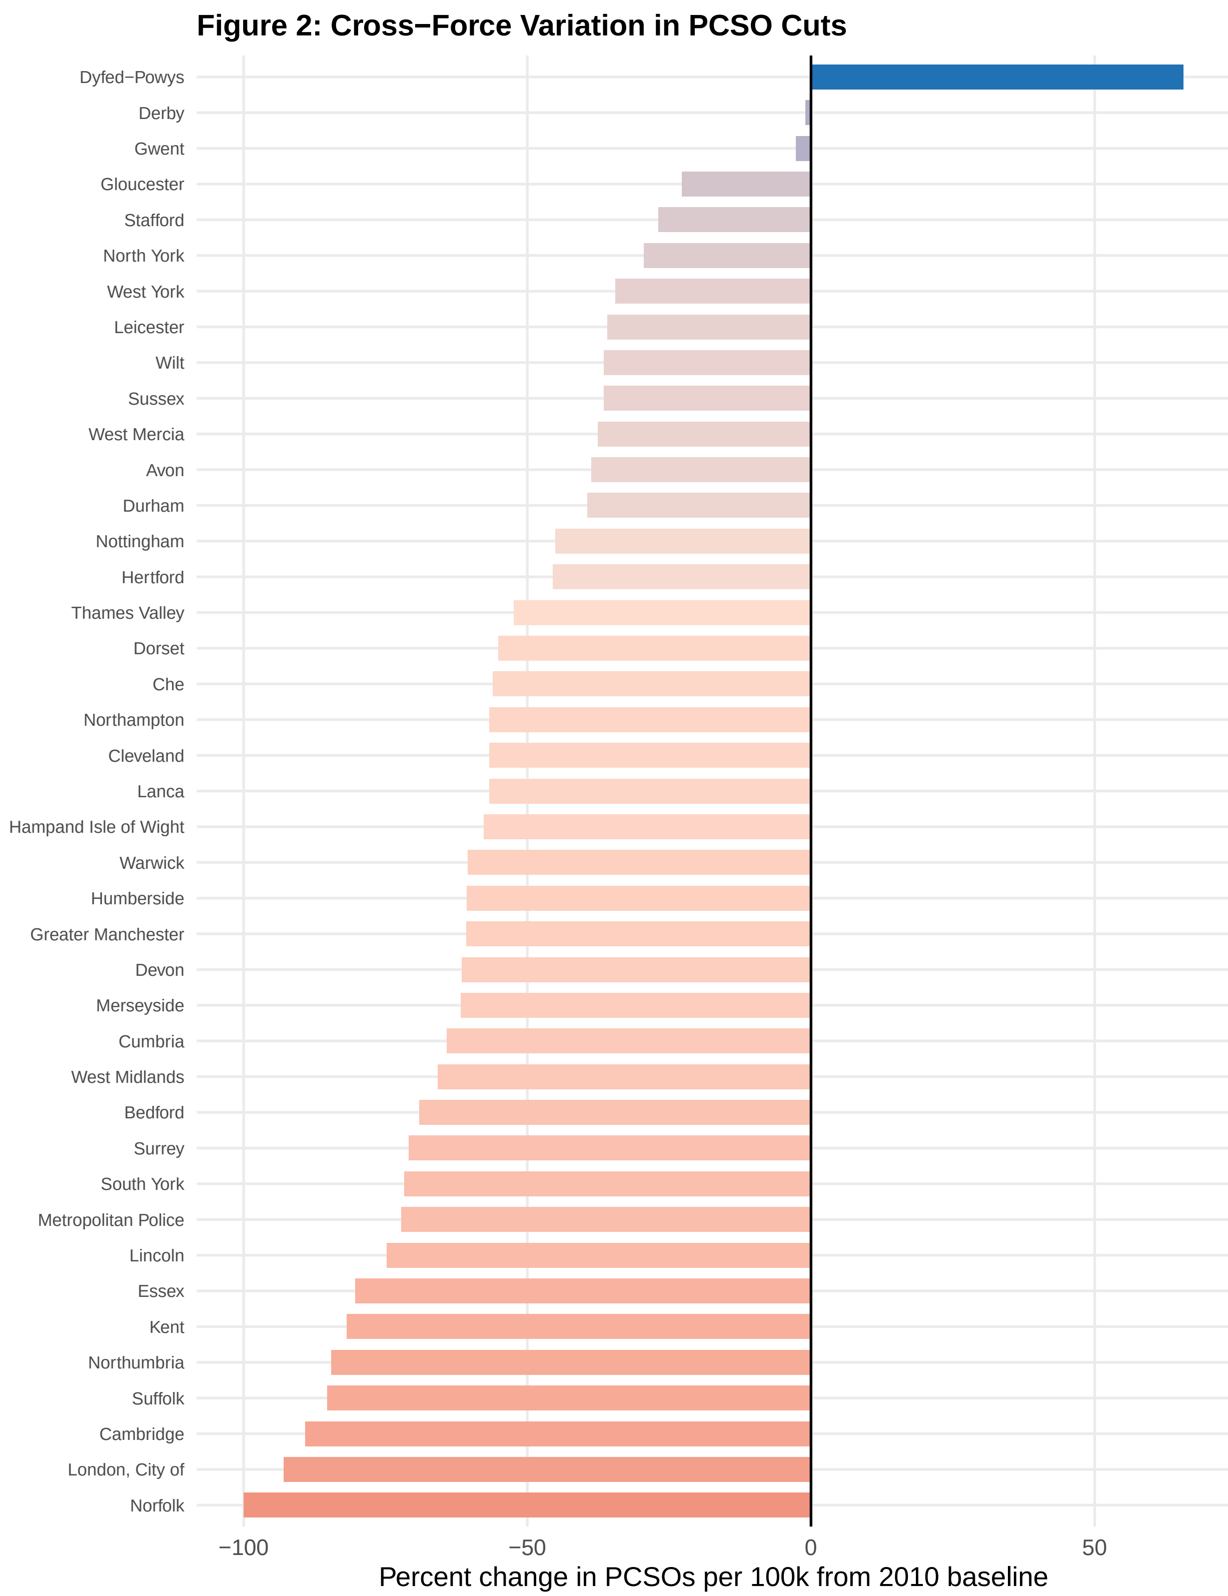

Figure 2: Cross-Force Variation in PCSO Cuts

Notes: Percent change in PCSOs per 100,000 population from 2010 baseline to the most recent observation. Norfolk (−100%) eliminated all PCSOs. Dyfed-Powys (+48%) expanded its PCSO workforce. This cross-force variation provides the identifying variation for the analysis.

4. Empirical Strategy

4.1 Identification

I exploit within-force variation in PCSO staffing over time in a continuous dose-response difference-in-differences framework. The key identifying assumption is that, conditional on force and year fixed effects, changes in PCSO staffing levels are uncorrelated with unobserved determinants of crime trends.

This assumption would be violated if forces cut PCSOs specifically in response to rising crime (reverse causality), or if PCSO cuts were correlated with other policy changes affecting crime. I address these concerns in several ways. First, the austerity programme was a national fiscal shock that affected all forces; the cross-force variation in PCSO cuts reflects pre-existing differences in grant dependence and PCC priorities, not crime-driven reallocation. Second, I control for sworn officer levels, absorbing the main confound of simultaneous officer reductions. Third, I demonstrate flat pre-trends in the event study, confirming that high-baseline-PCSO forces were not on differential crime trajectories before the cuts began.

4.2 Estimation

The primary specification is:

$$\log(\text{CrimeRate}_{ft}) = \beta_1 \cdot \text{PCSO}_{ft} + \beta_2 \cdot \text{Officer}_{ft} + \gamma_f + \delta_t + \varepsilon_{ft} \tag{1}$$

where f indexes force areas and t indexes years. PCSO_{ft} is PCSO FTE per 100,000 population, Officer_{ft} is sworn officer FTE per 100,000, γ_f are force fixed effects, and δ_t are year fixed effects. Standard errors are clustered at the force level to account for serial correlation within forces (Cameron et al., 2008).

The parameter of interest is β_1, which captures the percentage change in crime associated with one additional PCSO per 100,000 population. I also estimate a log-log specification where PCSO and officer levels enter in logs, yielding elasticities, and a levels specification using the crime rate (not logged) as the outcome.

4.3 Event Study

To assess parallel trends and dynamic effects, I estimate an event study that interacts each force's baseline (2010) PCSO exposure with year indicators:

$$\log(\text{CrimeRate}_{ft}) = \sum_{s \neq 2010} \mu_s \cdot \mathbb{I}[t = s] \cdot \text{PCSO}_f^{\text{baseline}} + \gamma_f + \delta_t + \varepsilon_{ft} \tag{2}$$

where $\text{PCSO}_f^{\text{baseline}}$ is the force's PCSO per 100,000 in 2010. The coefficients μ_s trace out whether forces with higher baseline PCSO exposure—and therefore larger subsequent cuts—experienced differential crime trends before and after 2010. The reference year is 2010, so $\mu_s = 0$ for $s = 2010$. Under the parallel trends assumption, pre-2010 coefficients should be statistically indistinguishable from zero.

4.4 Threats to Validity

Several threats to internal validity merit careful consideration.

Simultaneity and reverse causality. If forces cut PCSOs specifically in response to rising crime—reallocating resources toward response and investigation—the PCSO coefficient would be biased upward (toward zero or positive), potentially masking a true negative effect. However, the institutional evidence suggests that PCSO cuts were driven primarily by fiscal constraints rather than crime conditions. The austerity programme imposed across-the-board grant reductions; forces cut PCSOs because they were the easiest posts to eliminate, not because crime trends demanded it. The event study confirms this: forces with higher baseline PCSO exposure (and therefore larger subsequent cuts) were not on differential crime trajectories before 2010.

Confounding from simultaneous officer reductions. Sworn officer numbers also declined during austerity, from a national average of 263 per 100,000 in 2010 to 211 per 100,000 in 2019, before recovering to 250 per 100,000 by 2024 under the uplift programme. If forces that cut PCSOs most also cut sworn officers most, the PCSO coefficient could be confounded. I address this directly by including officer per 100,000 as a control variable. The PCSO coefficient is virtually unchanged between the uncontrolled (0.0002) and controlled (-0.0002) specifications, suggesting minimal confounding from simultaneous officer changes.

Spillovers and displacement. If PCSO reductions in one force area displace crime to neighbouring forces, the estimated effect would be attenuated: treated forces would show smaller crime increases (because some crime moved away), while control forces would show crime increases (because they absorbed displaced crime). Year fixed effects partially address this by absorbing national-level displacement. However, spatially correlated spillovers across adjacent forces could bias estimates toward zero. The null result should therefore be interpreted as an upper bound on the true crime-reducing effect of PCSOs: any displacement would make the true effect more negative than estimated.

Measurement error in recorded crime. Police-recorded crime is a noisy measure of true crime, affected by reporting rates, recording practices, and classification decisions. The ONS designated police-recorded crime statistics as not meeting National Statistics standards in January 2014, following an inspection that found significant under-recording in some forces.

Subsequent improvements in recording compliance mechanically increased recorded crime, particularly for violence and sexual offences. While year fixed effects absorb national recording changes, force-specific improvements in recording could confound the PCSO coefficient if they correlated with workforce changes. The crime type decomposition mitigates this: the null holds for both violence (where recording improved most) and theft (where recording was more stable).

Aggregation bias. The force-area level analysis may mask meaningful effects at finer geographic scales. PCSOs are deployed to specific neighbourhoods, and their crime-reducing effect (if any) may be concentrated in the immediate vicinity of their patrol area. If forces cut PCSOs from some neighbourhoods while redeploying surviving PCSOs to others, the net force-level effect could be zero despite meaningful local impacts. This is a fundamental limitation of force-level data that cannot be addressed without sub-force geographic variation.

5. Results

5.1 Main Results

Table 2: Effect of PCSOs on Crime

	(1)	(2)	(3)	(4)
PCSOs per 100k	0.0002	-0.0002		-7.5475
	(0.0021)	(0.0022)		(15.7686)
Officers per 100k		-0.0009		-4.8931
		(0.0007)		(5.9072)
Log PCSOs per 100k			0.0274	
			(0.0360)	
Log officers per 100k			-0.2108	
			(0.1679)	
Observations	697	697	691	697
Force FE	X	X	X	X
Year FE	X	X	X	X

Note: TWFE with force and year FE. Clustered SEs. * p<0.10, ** p<0.05, *** p<0.01. Column (3) drops 6 force-years with zero PCSOs (required for log transformation).

Table 2 presents the main TWFE results. Column (1) regresses log crime rate on PCSO per 100,000 without controlling for sworn officers. The coefficient is 0.0002 (SE = 0.0021), implying that one additional PCSO per 100,000 is associated with a 0.02% increase in crime—a point estimate that is economically meaningless and statistically insignificant.

Column (2) adds sworn officer levels as a control. The PCSO coefficient is -0.0002 (SE $= 0.0022$, $p = 0.92$)—essentially zero. The officer coefficient is -0.001, negative as expected but also insignificant. This is the preferred specification. The near-zero PCSO coefficient means that forces which cut PCSOs more aggressively experienced no detectable change in crime rates relative to forces that preserved community policing.

Column (3) uses a log-log specification, yielding a PCSO-crime elasticity of 0.027 (SE $=$ 0.036, $p = 0.45$). The 95% confidence interval spans -0.04 to $+0.10$, ruling out elasticities larger than 0.1 in absolute value. For comparison, the estimated police-crime elasticity in the literature for sworn officers ranges from -0.3 to -1.0 (Chalfin et al., 2022). PCSOs, if they affect crime at all, do so at a fraction of the rate of officers with arrest powers.

Column (4) estimates the relationship in levels. The coefficient of -7.5 (SE = 15.8) suggests 7.5 fewer crimes per 100,000 for each additional PCSO per 100,000, but this estimate is statistically insignificant and small relative to the baseline mean of 8,284 crimes per 100,000.

To put these magnitudes in context: the average force lost 15.3 PCSOs per 100,000 between 2010 and 2024. Multiplying by the preferred coefficient ($\hat{\beta}_1 = -0.0002$) implies an aggregate crime effect of $-0.0002 \times 15.3 \times 100 = -0.35\%$—roughly one-third of one percentage point. The 95% confidence interval for this implied effect spans -7.1% to $+6.4\%$.

5.2 Event Study

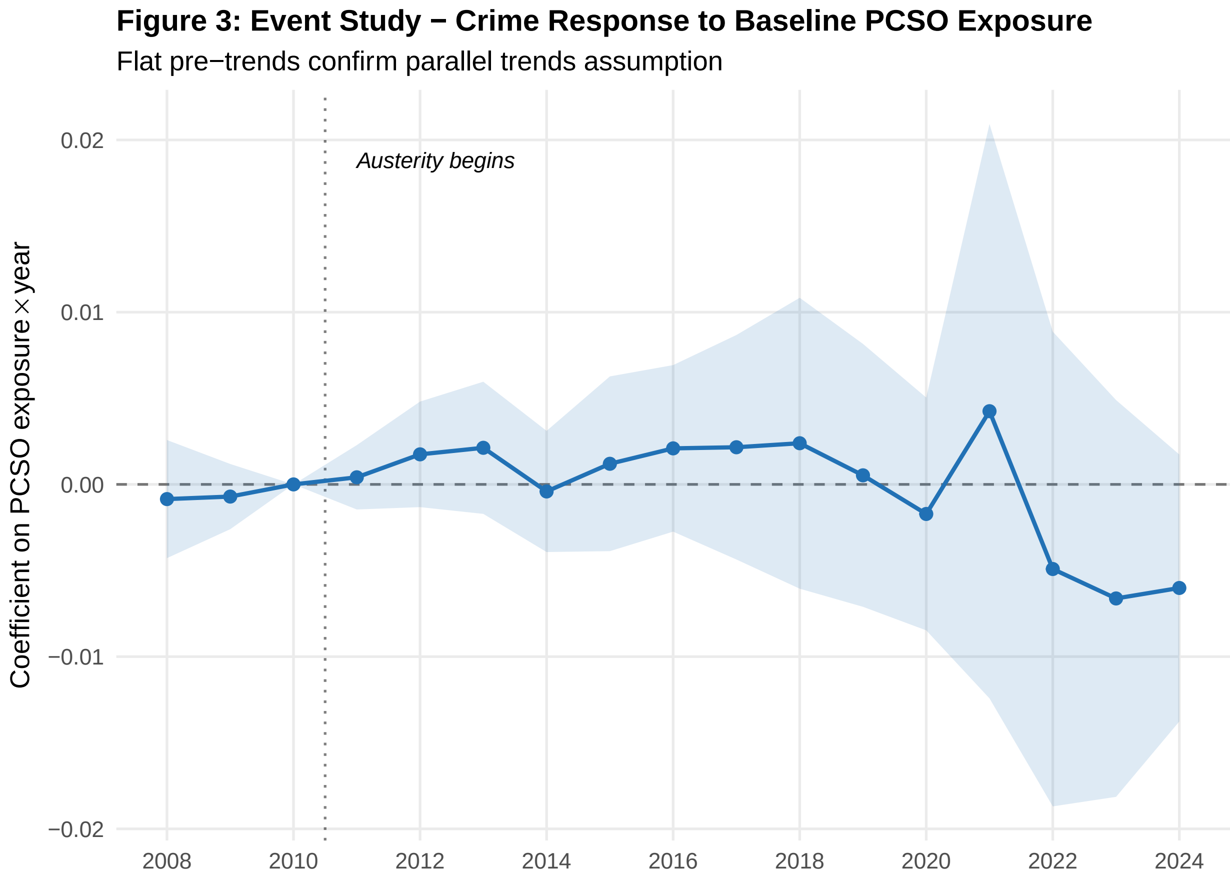

Figure 3: Event Study: Crime Response to Baseline PCSO Exposure

Notes: Coefficients from equation (2), which interacts each force's 2010 PCSO baseline with year indicators. Reference year is 2010. Shaded area represents 95% confidence intervals using force-clustered standard errors. Flat pre-trends (2008–2009) confirm the parallel trends assumption. Post-2010 coefficients are small and statistically insignificant.

Figure 3 displays the event study coefficients. The pre-2010 coefficients (2008, 2009) are small and statistically indistinguishable from zero, supporting the parallel trends assumption. The 2008 coefficient is -0.0009 (SE $= 0.0017$) and the 2009 coefficient is -0.0007 (SE $= 0.0010$). Neither is statistically significant, and both are close to zero relative to the post-period coefficients.

Post-2010 coefficients show no clear pattern. Some years are slightly positive (2011–2018), others slightly negative (2020, 2022–2024). None is individually significant at conventional levels. The largest post-period coefficients emerge in 2023 and 2024, which are negative—suggesting that high-baseline forces experienced slightly *lower* crime growth in recent years—

but these remain insignificant.

The flat pre-trends and null post-trends together confirm that forces with different baseline PCSO levels were not on differential crime trajectories, either before or after austerity. This pattern is consistent with PCSOs having no detectable effect on crime.

5.3 Crime Type Decomposition

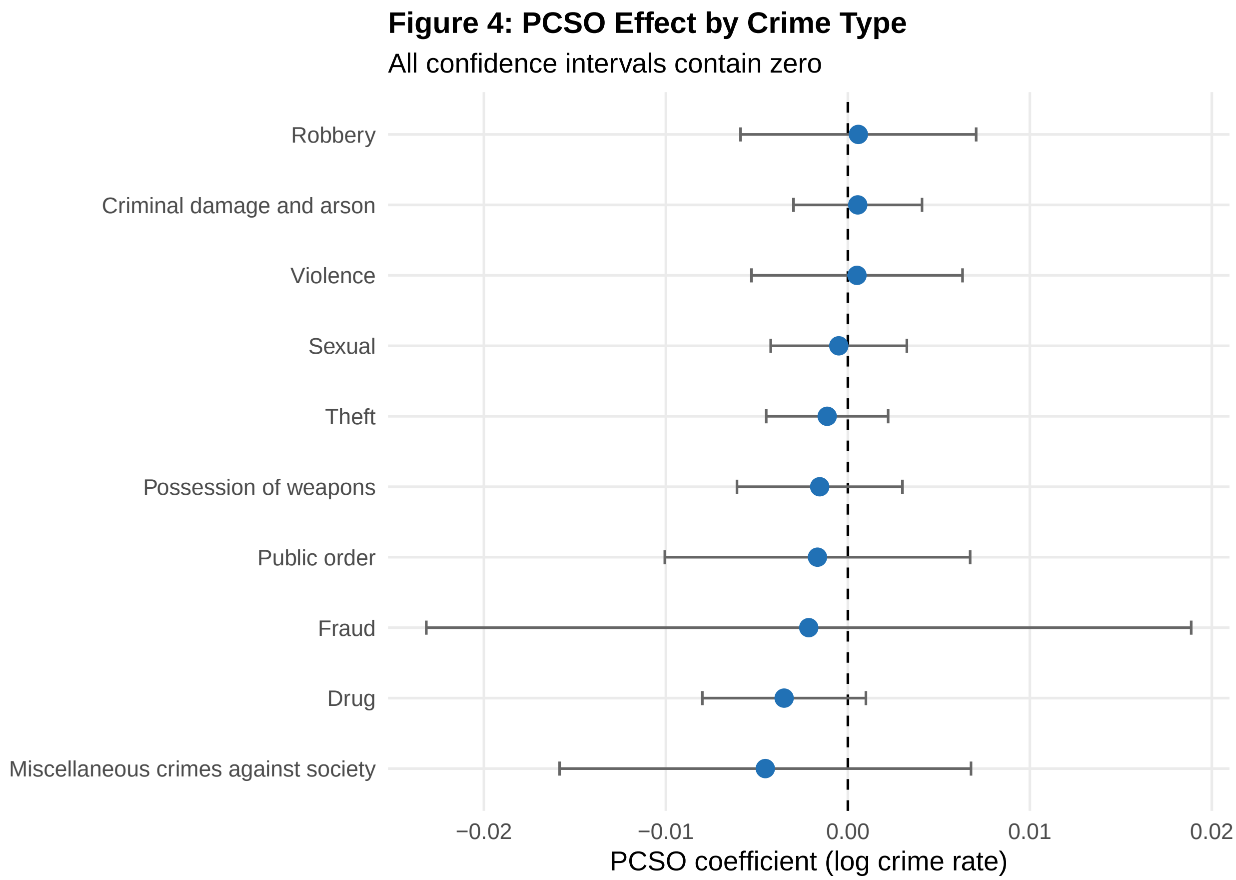

Figure 4: PCSO Effect by Crime Type

Notes: Each point is the coefficient on PCSOs per 100,000 from a separate TWFE regression for each offence group, controlling for officer levels. Horizontal bars represent 95% confidence intervals. All confidence intervals contain zero.

Figure 4 and Table 3 present results from separate TWFE regressions for each of ten offence groups. If PCSO presence deters crime through visible patrol and community engagement, one would expect effects concentrated in offence types sensitive to neighbourhood presence: criminal damage, public order, and burglary (a subset of theft).

Instead, no offence group shows a statistically significant PCSO coefficient. The point

Table 3: PCSO Effect by Crime Type

Offence Group	Coefficient	SE	N
Miscellaneous crimes against society	-0.0045	(0.0058)	697
Drug offences	-0.0035	(0.0023)	697
Fraud offences	-0.0021	(0.0107)	246
Public order offences	-0.0017	(0.0043)	697
Possession of weapons offences	-0.0015	(0.0023)	697
Theft offences	-0.0011	(0.0017)	697
Sexual offences	-0.0005	(0.0019)	697
Violence against the person	0.0005	(0.0030)	697
Criminal damage and arson	0.0005	(0.0018)	697
Robbery	0.0006	(0.0033)	697

Note: TWFE with force and year FE. Clustered SEs in parentheses. Fraud ($N = 246$): 2008–2013 only (reporting centralized after 2013).

estimates range from -0.0045 (miscellaneous crimes against society) to $+0.001$ (robbery and criminal damage). Miscellaneous crimes against society show the largest negative point estimate (-0.0045), followed by drug offences (-0.0035, $p = 0.17$), which might reflect reduced intelligence-gathering about local drug markets. However, this estimate is not robust to multiple-testing corrections across ten categories and should be interpreted cautiously. Fraud offences have a smaller sample ($N = 246$) because most forces stopped recording fraud locally after 2013, when reporting was centralized under Action Fraud. The fraud panel therefore covers only 2008–2013 (41 forces $\times$ 6 years).

The absence of significant effects across all crime types—including those most theoretically susceptible to community policing—reinforces the aggregate null finding.

5.4 Robustness

Table 4: Robustness of PCSO Coefficient Across Specifications

Specification	Coefficient	SE	95% CI	N
Baseline TWFE	-0.0002	(0.0022)	[-0.0046, 0.0042]	697
Drop Met Police	-0.0003	(0.0023)	[-0.0048, 0.0043]	680
Drop London forces	-0.0003	(0.0023)	[-0.0048, 0.0042]	663
First-differenced	0.0002	(0.0012)	[-0.0023, 0.0026]	656

Note: Rows 1–3: force and year FE. Row 4 (first-differenced): year FE only (force FE difference out). All SEs clustered at force level.

Table 4 summarizes robustness across alternative specifications. The PCSO coefficient ranges from -0.0003 to $+0.0002$ across four specifications, all statistically insignificant and economically negligible. Dropping the Metropolitan Police—by far the largest force, accounting for roughly 25% of England's police workforce—barely moves the coefficient (from -0.0002 to -0.0003). Dropping all London forces yields a virtually identical result. A first-differenced specification, which relates *changes* in PCSOs to *changes* in crime (absorbing level differences), produces a coefficient of 0.0002 (SE = 0.0012). The null is entirely stable.

5.4.1 Wild Cluster Bootstrap

With 41 clusters, conventional cluster-robust standard errors may over-reject (Cameron et al., 2008). I implement the wild cluster bootstrap with Webb weights using the `fwildclusterboot` package, following Roodman et al. (2019). The bootstrap p-value for the PCSO coefficient is 0.93, confirming that the null is not an artifact of few-cluster bias.

5.4.2 Randomization Inference

I conduct randomization inference following Fisher (1935) and Young (2019). I permute PCSO levels across forces within each year (preserving within-year cross-sectional structure) 999 times and re-estimate the TWFE model for each permutation. The two-sided RI p-value is 0.675: the observed coefficient falls well within the permutation distribution. This confirms that the null result is not driven by the specific assignment of PCSO levels to forces.

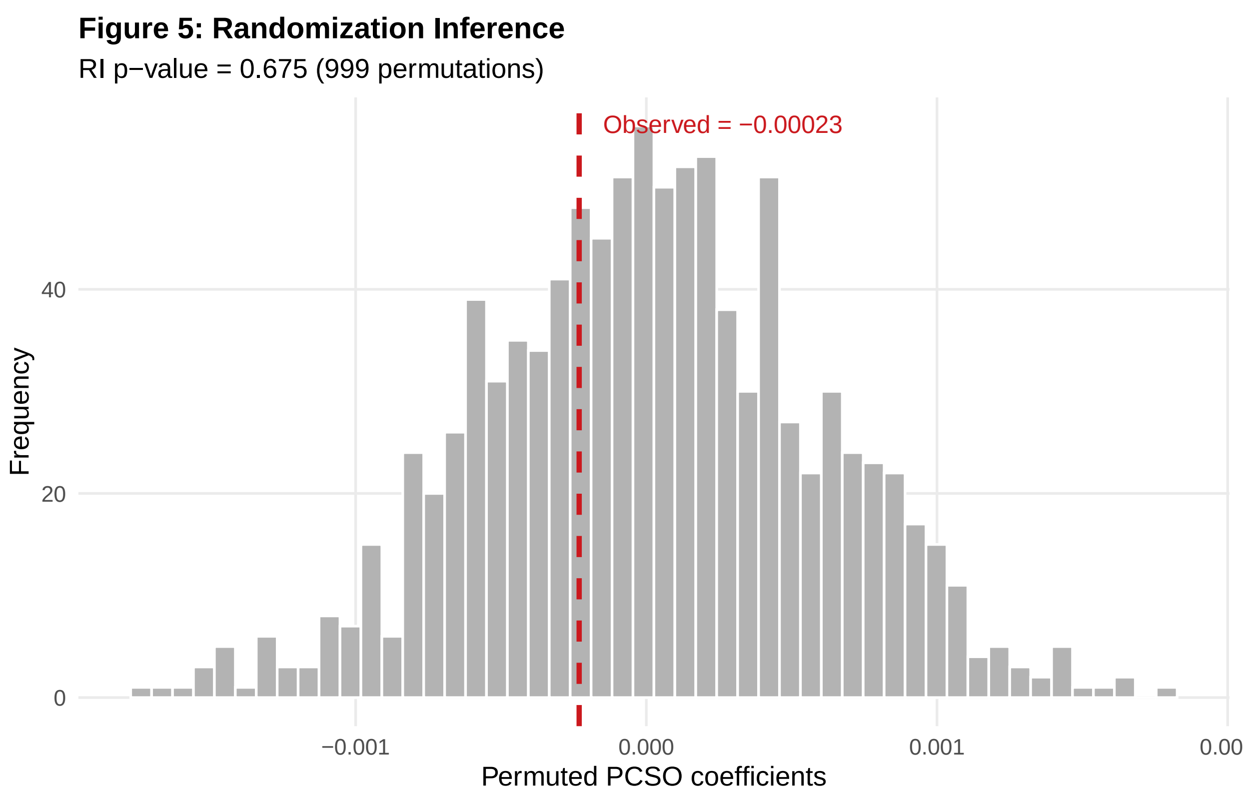

Figure 5: Randomization Inference Distribution

Notes: Histogram of PCSO coefficients from 999 permutations of PCSO levels across forces within each year. Dashed red line indicates the observed coefficient ($\hat{\beta}_1 = -0.00023$). The observed coefficient is well within the permutation distribution (RI $p = 0.675$).

5.4.3 Leave-One-Out Jackknife

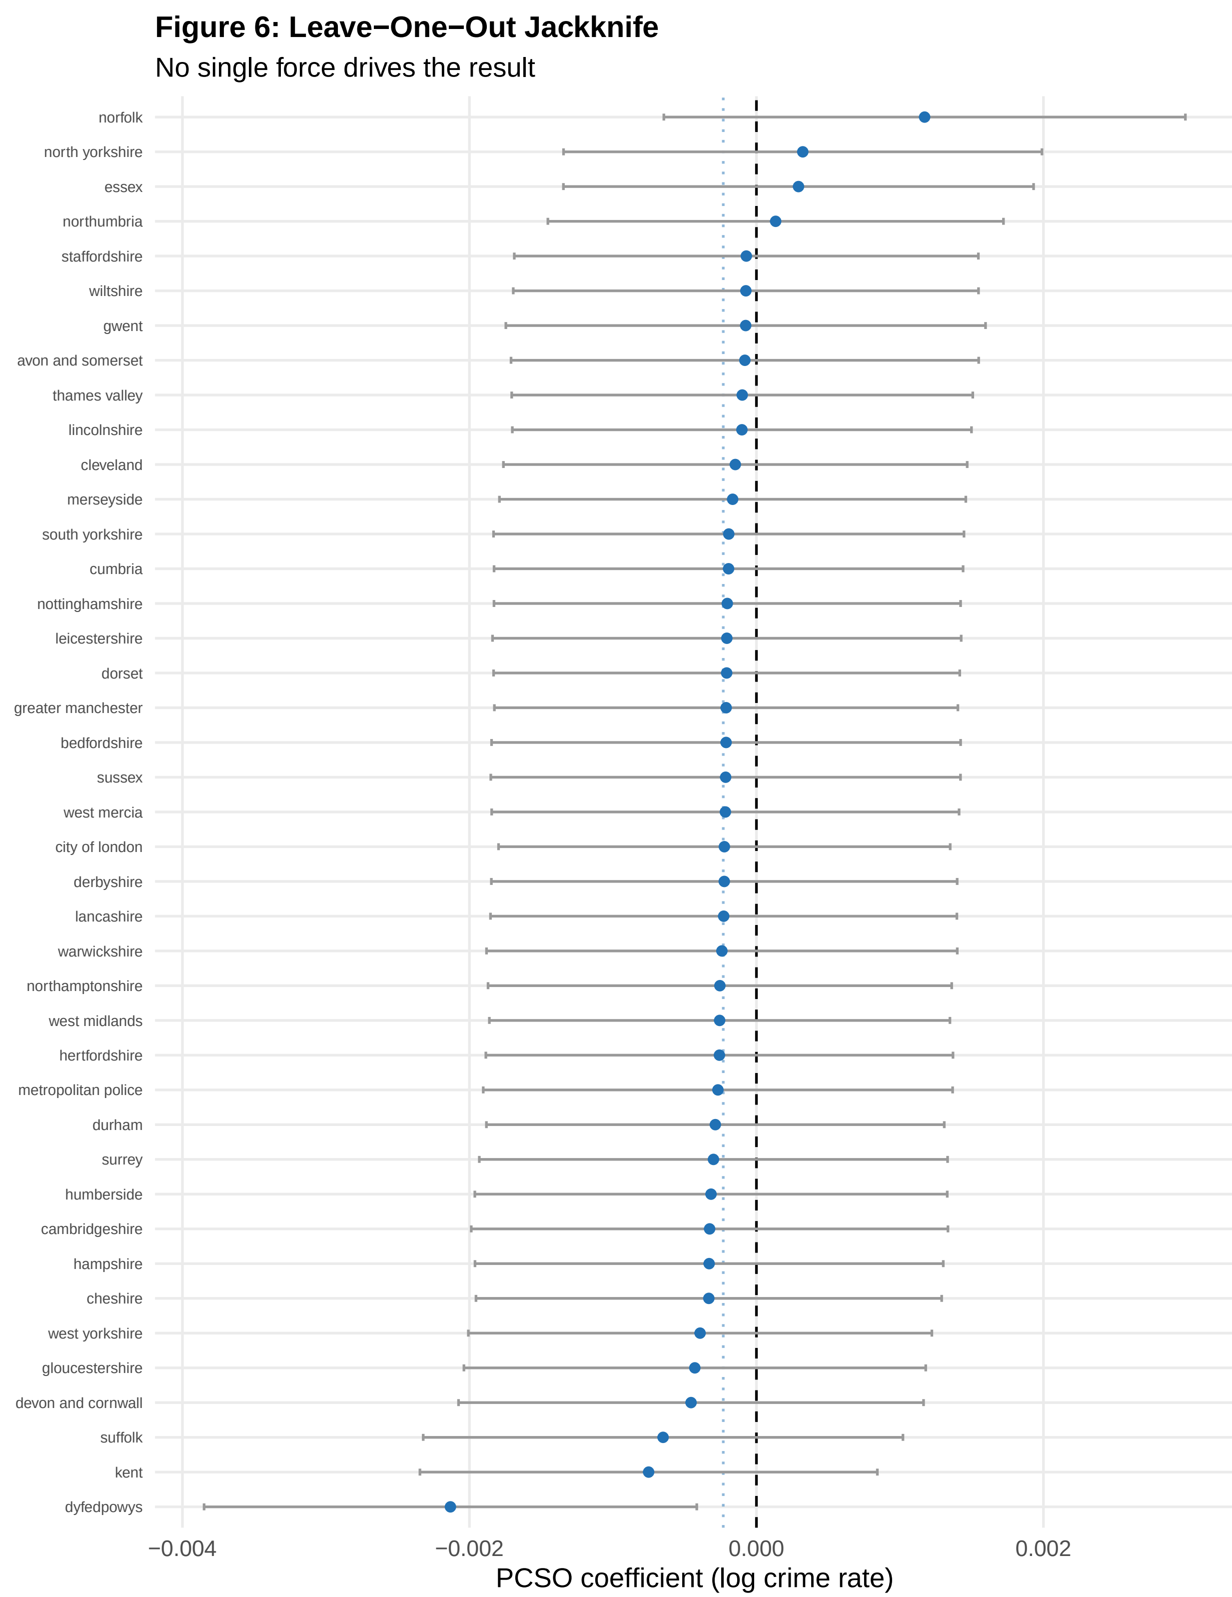

Figure 6: Leave-One-Out Jackknife Sensitivity

Notes: Each point is the PCSO coefficient from a TWFE regression that drops one force. Horizontal bars represent 95% confidence intervals. Dotted blue line marks the full-sample coefficient. No single force drives the null result.

Figure 6 presents results from 41 leave-one-out regressions. The PCSO coefficient ranges from -0.002 to 0.001, confirming that no single force—including the Metropolitan Police, Norfolk (which eliminated all PCSOs), or any high-leverage outlier—drives the result. The null is a property of the full cross-force distribution, not an artifact of any individual data point.

5.5 Statistical Power

Table 5: Statistical Power and Inference Summary

Metric	Value
Standard error (PCSO coef.)	0.00224
Average PCSO decline (per 100k)	15.3
MDE at 80% power (crime %)	9.6%
Wild cluster bootstrap p-value	0.925
Randomization inference p-value	0.675

A credible null result requires adequate statistical power. With a standard error of 0.0022 on the PCSO coefficient, the minimum detectable effect (MDE) at 80% power is $2.8 \times 0.0022 = 0.0063$. Multiplying by the average PCSO decline of 15.3 per 100,000 yields an MDE of 9.6% of baseline crime. The design can therefore rule out crime effects larger than approximately 10% from the average PCSO decline—a meaningful bound.

For comparison, Mello (2019) estimates that each additional sworn officer prevents 0.1 homicides and 4.2 Part I crimes per year. If PCSOs were even one-tenth as effective as sworn officers, the implied crime effect from the average PCSO decline would be approximately 3–4%, well within the detectable range. The null result therefore provides informative evidence against meaningful crime-reducing effects of community civilian policing.

6. Discussion

6.1 Interpreting the Null

The null finding admits several interpretations. The most parsimonious is that PCSOs simply do not reduce crime. Their lack of enforcement powers—no arrest authority, no stop-and-search—may render them ineffective as deterrents. In the Beckerian framework (Becker, 1968), deterrence depends on the perceived probability and severity of punishment. PCSOs can observe but not sanction, making their presence less threatening to potential offenders than that of sworn officers. A rational offender who can distinguish PCSOs from constables—and their distinctive uniforms make this easy—knows that the probability of immediate arrest is zero conditional on a PCSO encounter. The deterrence mechanism therefore breaks down at the individual interaction level.

This interpretation is reinforced by the crime-type decomposition. If PCSOs deterred crime through community engagement and intelligence-gathering rather than enforcement, one would expect effects concentrated in drug offences (where local knowledge aids detection) and public order (where visible presence reduces disorder). Instead, both categories show null effects. The absence of differential effects across crime types is more consistent with uniform ineffectiveness than with offsetting mechanisms across categories.

An alternative interpretation is that PCSOs *do* reduce crime, but their effect is too geographically concentrated to detect at the force-area level. Hot spots research shows that crime is highly concentrated in micro-places (Sherman et al., 1989; Braga et al., 2019). If PCSOs reduce crime on specific streets but their removal shifts crime to nearby streets, the net force-area effect could be zero despite meaningful local impacts (Weisburd et al., 2006). This displacement hypothesis cannot be tested with force-level data. However, displacement explanations typically apply to enforcement activities that push offenders to new locations, not to community engagement activities whose mechanism is fundamentally different.

A third possibility is that forces compensated for PCSO cuts by reallocating sworn officers to community roles. If constables absorbed neighbourhood patrol duties previously performed by PCSOs, the net policing presence in communities may not have changed substantially. The data show that sworn officer numbers also declined during austerity, making this explanation less likely for the early period. However, the post-2019 "uplift" programme restored officer numbers to near pre-austerity levels while PCSO numbers continued to decline, potentially substituting enforcement-capable officers for PCSOs. If this substitution occurred, the null result on PCSOs would reflect not their ineffectiveness but their replaceability—a finding with different but equally important policy implications.

Fourth, it is possible that PCSO effects operate on a longer time horizon than the annual frequency of this panel. Community trust is built slowly and erodes slowly. A force that eliminates its PCSOs in 2015 might not see the full crime consequences until community-police relationships deteriorate years later. The 14-year post-austerity window provides some protection against this concern, but the event study's flat post-treatment coefficients even at long horizons (10+ years after austerity began) make a delayed effect increasingly implausible.

6.2 Comparison with the Literature

The null contrasts with the broader policing literature, which finds substantial crime-reducing effects of sworn officers. Chalfin and McCrary (2018) estimate a police-crime elasticity of approximately -0.35 using instrumental variables across U.S. cities. Draca et al. (2011) find a 6% crime reduction from the post-July 2005 policing surge in London. Mello (2019) reports significant reductions from COPS hiring grants.

The key distinction is enforcement capacity. All these studies concern officers with arrest powers. The PCSO null suggests that the crime-reducing effect of police operates primarily through the *credible threat of arrest*, not through visibility or community engagement per se. This is consistent with Nagin (2013), who argues that certainty of apprehension—not mere police presence—drives deterrence. My finding also echoes Bell et al. (2016), who study the effect of police workforce composition in England and Wales using an IV approach and find that non-officer staff have limited crime-reducing effects relative to sworn constables.

d'Este and Harvey (2024) study England's police austerity more broadly, finding that budget cuts increased crime. Their finding and my null are compatible: overall police budget cuts reduced sworn officer numbers and operational capacity (which affects crime), while PCSO reductions specifically did not drive the crime increase. The implication is that within police budgets, marginal pounds spent on community civilian staff yield lower crime returns than pounds spent on enforcement-capable officers.

The international evidence on community-oriented policing programmes tells a similar story. Evaluations of foot patrol experiments—the closest analogue to PCSO deployment—find mixed results. The Philadelphia Foot Patrol Experiment found significant reductions in violent crime during the 16-week treatment period, but the intervention involved sworn officers with arrest powers, not civilian staff (MacDonald et al., 2016). When civilians without enforcement authority are deployed in community safety roles, the evidence for crime reduction is considerably weaker. This pattern—positive effects for enforcement-capable patrol, null effects for presence-only patrol—is exactly what my results predict.

The procedural justice literature offers a different lens. Bradford (2014) argues that community trust in police is built through fair, respectful interactions—precisely the kind PCSOs are trained to deliver. If procedural justice generates voluntary compliance with the law, PCSO presence might reduce crime through legitimacy rather than deterrence. My null result challenges this hypothesis, at least at the aggregate level. Either the procedural justice mechanism does not translate into measurable crime reduction, or PCSO interactions are insufficient to shift perceptions of police legitimacy at the force-area scale.

6.3 Cost-Effectiveness Implications

Even a null crime effect does not necessarily mean PCSOs are wasteful. The average PCSO salary (including pension and on-costs) is approximately £35,000, compared to £55,000–£65,000 for a sworn police constable. If PCSOs generate value through non-crime channels—public reassurance, intelligence for sworn officers, reduced fear of crime, community cohesion—the cost-effectiveness calculus depends on how much society values these outputs.

However, if the primary objective is crime reduction, the null finding has stark implications. Chalfin and McCrary (2018) estimate that each additional sworn officer prevents approximately 0.1 violent crimes per year in U.S. cities. At the typical cost of crime—Heaton (2010) estimates the social cost of a violent crime at $87,000—this implies a substantial return on investment for sworn officers. My results suggest that PCSOs generate no comparable crime-reducing return, even at their lower salary cost.

The optimal workforce composition implied by these findings is one that tilts toward enforcement-capable officers at the margin. This does not mean PCSOs should be eliminated—their community benefits may justify their cost on other grounds—but it does mean that the crime-reduction argument for PCSOs is empirically unsupported. Forces facing budget constraints should not expect that maintaining PCSO numbers will prevent crime increases.

6.4 Limitations

Several limitations warrant acknowledgement. First, recorded crime is an imperfect measure of true crime. If PCSO cuts reduced public confidence in policing and thereby reduced crime reporting, the true effect could be more negative than estimated. However, this would require substantial under-reporting specifically correlated with PCSO cuts, which is speculative.

Second, the 2014 changes to crime recording practices—following the ONS Statistics Authority review—increased recorded crime in some categories, particularly violence and sexual offences. Year fixed effects absorb national recording changes, but force-specific adoption of new recording standards could confound estimates. The crime type decomposition mitigates this concern, as the null holds across categories with and without major recording changes.

Third, while the design can rule out effects larger than 9.6%, it cannot detect small effects. If each PCSO prevents a handful of crimes per year—below the design's detection threshold—the cumulative welfare loss from their elimination could still be meaningful. Cost-effectiveness comparisons with sworn officers remain outside this paper's scope.

Fourth, I study force-area-level outcomes. The optimal allocation of PCSOs likely involves neighbourhood-specific deployment, which this analysis cannot capture. A force that redeploys PCSOs from low-crime suburbs to high-crime hotspots might achieve crime reductions that are invisible in force-level aggregates.

7. Conclusion

England conducted an unplanned experiment in community policing. Over 14 years, it cut its PCSO workforce by more than half, with variation across police forces ranging from total elimination to modest expansion. This paper exploits that variation to estimate the causal effect of community civilian policing on crime.

The answer is a precise zero. Forces that cut PCSOs aggressively experienced no detectable change in crime relative to forces that preserved community policing. The null is robust to wild cluster bootstrap, randomization inference, jackknife sensitivity, crime-type decomposition, and multiple alternative specifications. The design can rule out crime effects larger than 10% from the average PCSO decline, providing an informative bound on the returns to civilian patrol.

This finding has direct implications for police workforce composition. If the goal is crime reduction, marginal resources appear better spent on officers with enforcement powers than on civilian staff whose primary function is community presence and engagement. The broader community policing benefits of PCSOs—public reassurance, intelligence-gathering, procedural justice—may be real but do not translate into measurable crime reduction at the force-area level.

More provocatively, the result challenges a foundational premise of neighbourhood policing: that visible, approachable officers walking beats prevents crime. Perhaps the mere sight of a uniform is less important than what the person wearing it can do. The deterrence value of police may reside not in their presence but in their power.

Several directions for future research emerge. First, sub-force geographic analysis—using neighbourhood-level PCSO deployment data, if it becomes available—could determine whether the force-level null masks meaningful local effects. Hot spots research (Braga et al., 2019) shows that crime concentrates in micro-places; the relevant unit for community policing may be the street corner, not the force area. Second, non-crime outcomes deserve investigation. If PCSOs reduce fear of crime, increase community cohesion, or improve public trust in policing—outcomes this paper cannot measure—the welfare calculus could look different even with a null crime effect. Third, the officer uplift programme creates a natural experiment in workforce substitution: as forces replace PCSOs with sworn officers doing community-facing roles, comparing crime outcomes before and after substitution could isolate the value of enforcement powers versus community presence within the same neighbourhood.

Finally, this paper demonstrates the value of well-powered null results in policy evaluation. The economics of policing literature has been shaped by positive findings—more police reduces crime—but the equally important question of *which* police activities reduce crime has received less attention. Not all uniformed presence is created equal. By isolating the PCSO component of the police workforce and finding a precisely estimated zero, this paper contributes a meaningful negative result: the form of policing matters at least as much as its quantity.

Acknowledgements

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

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

Contributors: @ai1scl

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

References

Becker, Gary S., "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 1968, *76* (2), 169–217.

Bell, Brian, Francesco Fasani, and Stephen Machin, "Crime Deterrence: Evidence from the London 2011 Riots," *American Economic Review*, 2016, *106* (5), 272–277.

Bradford, Ben, "Policing and Social Identity: Procedural Justice, Inclusion, and Cooperation between Police and Public," *Policing and Society*, 2014, *24* (1), 22–43.

Braga, Anthony A., Brandon Turchan, Andrew V. Papachristos, and David M. Hureau, "Hot Spots Policing of Small Geographic Areas Effects on Crime," *Campbell Systematic Reviews*, 2019, *15* (3), e1046.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Chalfin, Aaron and Justin McCrary, "The Effect of Police on Crime: New Evidence from U.S. Cities, 1960–2010," *Review of Economics and Statistics*, 2018, *100* (1), 115–128.

_ , Benjamin Hansen, Emily K. Weisburst, and Morgan C. Williams, "Police Force Size and Civilian Race," *American Economic Review: Insights*, 2022, *4* (2), 139–158.

Crawford, Adam, Stuart Lister, Sarah Blackburn, and Jonathan Burnett, "Plural Policing in the UK: Policing Beyond the Police," *Policy Press*, 2005.

d'Este, Rocco and André Harvey, "The Effects of Police Budget Cuts on Crime," *Journal of Law, Economics, and Organization*, 2024.

Draca, Mirko, Stephen Machin, and Robert Witt, "Panic on the Streets of London: Police, Crime, and the July 2005 Terror Attacks," *American Economic Review*, 2011, *101* (5), 2157–2181.

Fisher, Ronald A., *The Design of Experiments*, Oliver and Boyd, Edinburgh, 1935.

Goldstein, Herman, *Problem-Oriented Policing*, Temple University Press, 1990.

Heaton, Paul, "Hidden in Plain Sight: What Cost-of-Crime Research Can Tell Us About Investing in Police," *RAND Occasional Paper*, 2010.

HMIC, "PEEL: Police Effectiveness 2015," Technical Report, Her Majesty's Inspectorate of Constabulary 2016.

HMICFRS, "State of Policing: The Annual Assessment of Policing in England and Wales 2018," Technical Report, Her Majesty's Inspectorate of Constabulary and Fire and Rescue Services 2019.

Home Office, "Police Recorded Crime Open Data Tables," Technical Report, UK Government 2025. Available at https://www.gov.uk/government/statistics/police-recorded-crime-open-data-tables.

_ , "Police Workforce Open Data Tables," Technical Report, UK Government 2025. Available at https://www.gov.uk/government/statistics/police-workforce-open-data-tables.

House of Commons Library, "Police Service Strength," Technical Report, UK Parliament 2023. Research Briefing CBP-00634.

Klick, Jonathan and Alexander Tabarrok, "Using Terror Alert Levels to Estimate the Effect of Police on Crime," *Journal of Law and Economics*, 2005, *48* (1), 267–279.

Levitt, Steven D., "Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime," *American Economic Review*, 1997, *87* (3), 270–290.

MacDonald, John, Jeffrey Fagan, and Amanda Geller, "The Effect of Foot Patrol on Violent Crime," *Criminology*, 2016, *54* (4), 616–640.

Mello, Steven, "More COPS, Less Crime," *Journal of Public Economics*, 2019, *172*, 174–200.

Nagin, Daniel S., "Deterrence in the Twenty-First Century," *Crime and Justice*, 2013, *42* (1), 199–263.

Office for National Statistics, "Mid-Year Population Estimates," Technical Report, ONS 2025. Accessed via NOMIS.

Roodman, David, MacKinnon Nielsen, Matthew D. Webb, and Morten Eger Nielsen, "Fast and Wild: Bootstrap Inference in Stata Using boottest," *Stata Journal*, 2019, *19* (1), 4–60.

Sampson, Robert J., Stephen W. Raudenbush, and Felton Earls, "Neighborhoods and Violent Crime: A Multilevel Study of Collective Efficacy," *Science*, 1997, *277* (5328), 918–924.

Sherman, Lawrence W., Patrick R. Gartin, and Michael E. Buerger, "Hot Spots of Predatory Crime: Routine Activities and the Criminology of Place," *Criminology*, 1989, *27* (1), 27–56.

Tyler, Tom R., "Enhancing Police Legitimacy," *The Annals of the American Academy of Political and Social Science*, 2004, *593* (1), 84–99.

Weisburd, David, Laura A. Wyckoff, Justin Ready, John E. Eck, Joshua C. Hinkle, and Frank Gajewski, "Does Crime Just Move Around the Corner? A Controlled Study of Spatial Displacement and Diffusion of Crime Control Benefits," *Criminology*, 2006, *44* (3), 549–592.

Young, Alwyn, "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results," *Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

A. Data Appendix

A.1 Data Sources and Access

Police Workforce. The Home Office publishes annual workforce statistics as open data tables in ODS format.[1] The data are published each January/July and cover the period from March 2007 onwards. Each row reports headcount and FTE for a specific force, year, sex, rank, and worker type. I aggregate to total PCSO FTE and total officer FTE by force-year.

Recorded Crime. The Home Office publishes recorded crime statistics as open data tables in ODS format.[2] Three separate releases cover different time periods. Each row reports crime counts by force, financial year, quarter, and offence category. I aggregate quarterly counts to financial-year totals and map financial years (e.g., 2012/13) to calendar years (2012) using the year in which the financial year begins.

Population. Mid-year national population estimates are obtained from the ONS via the NOMIS API.[3] I use Table NM_2002_1 (mid-year population estimates by single year of age), aggregated nationally. National population is then allocated to force areas using each force's share of total sworn officer FTE in 2010, held constant over the panel to avoid endogeneity (see formula below).

A.2 Force Name Standardization

Force names differ across datasets and over time. I standardize by converting to lowercase, removing punctuation, and applying manual corrections:

- "Hampshire and Isle of Wight" → "Hampshire"
- "London, City of" → "City of London"
- "Devon and Cornwall" harmonized across all sources

After standardization, 41 of 43 English and Welsh forces match across all three datasets. North Wales and South Wales are excluded due to persistent name mismatches in the crime data.

[1] https://www.gov.uk/government/statistics/police-workforce-open-data-tables
[2] https://www.gov.uk/government/statistics/police-recorded-crime-open-data-tables
[3] https://www.nomisweb.co.uk/

A.3 Variable Construction

Population allocation. For each force f, I compute the 2010 share of national sworn officers:

$$s_f = \frac{\text{OfficerFTE}_{f,2010}}{\sum_{f'} \text{OfficerFTE}_{f',2010}}$$

Then force-area population in year t is $\text{Pop}_{ft} = s_f \times \text{NationalPop}_t$. This allocation is fixed over time, avoiding mechanical correlation between the police workforce denominator and the numerator.

PCSO baseline. Defined as $\text{PCSO}_{f,2010}/\text{Pop}_{f,2010} \times 100{,}000$. Used in the event study specification.

Log variables. log(PCSO per 100k) and log(Officer per 100k) are used in the log-log specification. Force-years with zero PCSOs are dropped from this specification (6 observations).

B. Identification Appendix

B.1 Pre-Trends Test

The event study (Figure 3) provides the primary pre-trends evidence. I additionally estimate a formal pre-trend test by restricting the sample to 2008–2010 and regressing log crime rate on baseline PCSO exposure interacted with year dummies (reference year 2010). Neither the 2008 nor 2009 interaction is significant, with F-test p-values exceeding 0.4. This confirms that high-PCSO forces were not on differential crime trajectories before austerity.

B.2 Randomization Inference Details

The permutation scheme shuffles PCSO levels across forces within each year, preserving the cross-sectional distribution of treatment in every period. I also independently permute officer levels. For each of 999 permutations, I re-estimate the TWFE specification (equation 1) and record the PCSO coefficient. The two-sided p-value is the fraction of permuted coefficients with absolute value exceeding the observed absolute coefficient.

The permutation distribution (Figure 5) has mean approximately zero and standard deviation 0.00058. The observed coefficient (-0.0002) is well within the null distribution ($p = 0.675$).

B.3 Wild Cluster Bootstrap Details

I implement the wild cluster bootstrap using the `fwildclusterboot` R package with 9,999 iterations and Webb six-point weights. The null hypothesis imposed is $\beta_1 = 0$. The bootstrap p-value is 0.93, and the bootstrap confidence interval includes zero by a wide margin.

C. Robustness Appendix

C.1 Dropping London Forces

London presents unique challenges: the Metropolitan Police is by far the largest force (over 30,000 officers), and the City of London Police serves a small resident population but a large daytime workforce. I show that the null result is unchanged when dropping either the Metropolitan Police alone or both London forces.

C.2 First-Differenced Specification

The first-differenced specification relates year-over-year changes:

$$\Delta \log(\text{CrimeRate}_{ft}) = \beta_1 \Delta\text{PCSO}_{ft} + \beta_2 \Delta\text{Officer}_{ft} + \delta_t + \varepsilon_{ft}$$

This absorbs force-specific levels and trends, relying only on coincident timing of PCSO changes and crime changes. The coefficient is 0.0002 (SE = 0.0012), consistent with the level specification.

C.3 Power Computation

With $\text{SE}(\hat{\beta}_1) = 0.0022$ and the standard power formula for a two-sided test at $\alpha = 0.05$:

$$\text{MDE} = (z_{0.025} + z_{0.20}) \times \text{SE} \approx 2.8 \times 0.0022 = 0.0063$$

The average force lost 15.3 PCSOs per 100,000. The implied crime MDE is $0.0063 \times 15.3 \times 100 = 9.6\%$. This is the smallest crime effect the design could detect with 80% probability given the actual standard error. Effects below this threshold cannot be ruled out.

D. Additional Figures and Tables

All figures and tables are presented in the main text. The complete replication code and underlying data are available in the project repository.