

Forced Into the Light: Capital Controls and the Accidental Formalization of Greece’s Shadow Economy

APEP Autonomous Research* @ai1scl

March 11, 2026

Abstract

Can payment infrastructure shocks permanently reduce informality? I study Greece’s June 2015 capital controls, which capped ATM withdrawals at EUR 60/day while exempting card transactions—creating an accidental experiment in the relative cost of cash versus electronic payments. Cross-sector variation in pre-crisis cash dependence reveals perfect monotonic ordering: fuel retail (90% cash) fell 14.2%, food (75%) fell 7.0%, non-specialized (55%) fell 3.4%. A synthetic control comparison shows aggregate Greek retail turnover fell 10.4 index points below the counterfactual, with the gap *widening* after controls were lifted—consistent with hysteresis through sunk POS investment. Despite the turnover decline, Greece’s VAT revenue as a share of GDP rose 12.4 index points above the donor-country trend ($p = 0.008$), suggesting formalization partially offset the contractionary shock. Payment infrastructure, once imposed, can durably shrink the shadow economy.

JEL Codes: E26, H26, O17, E42, F38

Keywords: capital controls, shadow economy, informality, payment infrastructure, Greece, synthetic control, VAT compliance

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

On the evening of June 27, 2015, Greek Prime Minister Alexis Tsipras announced a national referendum on the terms of the country’s bailout. By the next morning, the European Central Bank had frozen Emergency Liquidity Assistance, and Greek banks were shuttered. When ATMs reopened on June 29, they dispensed at most EUR 60 per card per day. But card payments at point-of-sale terminals faced no cap. Overnight, the relative price of cash rose sharply—not through deliberate tax policy, but through the plumbing of crisis management.

This paper argues that Greece’s capital controls produced an accidental experiment in formalization. Greece entered the crisis with a shadow economy estimated at 22–27% of GDP (Medina and Schneider, 2018; Schneider et al., 2010), one of the largest in the EU. Much of this informality was sustained by cash: transactions off the books leave no electronic trail, and sectors with high cash intensity—fuel retail, food markets, small trade—had long been the locus of VAT evasion (European Commission, 2019). The capital controls did not target informality. They targeted capital flight. But by making cash scarce and card payments the path of least resistance, they forced a large share of previously informal transactions onto electronic rails.

To identify the effect, I triangulate three complementary empirical strategies: a cross-sector difference-in-differences exploiting cash-intensity variation, a synthetic control method quantifying the aggregate counterfactual, and a VAT-to-GDP analysis testing the formalization mechanism. The first strategy exploits within-Greece variation in pre-crisis cash dependence across retail subsectors. If capital controls merely reduced aggregate demand, all sectors should contract proportionally. If they additionally disrupted the cash-based shadow economy, high-cash sectors should suffer disproportionately—not because demand for fuel or food fell more, but because the portion of turnover that had been invisible to tax authorities was now being recorded. Treatment intensity is measured by the pre-2015 share of transactions conducted in cash, drawn from the ECB’s Survey on the Use of Cash by Households (European Commission, 2019).

The results are striking and monotonic. Fuel retail, where 90% of transactions were cash-based before 2015, experienced a 14.2% drop in the turnover index in July 2015. Food and beverages (75% cash) fell 7.0%. Non-specialized retail (55% cash) fell 3.4%. A cross-sector difference-in-differences specification confirms this ordering: a one-percentage-point increase in pre-2015 cash share is associated with a 0.62 index-point larger decline in reported turnover ($\hat{\beta} = -61.86$, $SE = 10.53$), with sector and time fixed effects. The binary specification—high-cash sectors ($\geq 75\%$) versus low—yields a 19.24 index-point gap. With only three sectors, cluster-robust inference is fragile (wild cluster bootstrap p -values are 0.29 and 0.16), but the

monotonic pattern itself—consistent with the theory and visible without any econometric adjustment—provides compelling design-based evidence.

I complement the cross-sector analysis with a synthetic control method (SCM) comparison of aggregate Greek retail turnover against a data-driven counterfactual (Abadie et al., 2010, 2015). The synthetic Greece, constructed from 14 EU donor countries, tracks actual Greece through the sovereign debt crisis and diverges after June 2015. The average gap over the full post-period is 10.43 index points, or approximately 9% of the pre-treatment mean. However, the aggregate SCM evidence is not conclusive on its own: Greece’s high pre-treatment volatility during the debt crisis produces a pre-treatment RMSPE that exceeds the post-treatment RMSPE, limiting the power of permutation-based inference. The value of the SCM lies in quantifying the aggregate magnitude and demonstrating persistence, not in delivering a single sharp test.

Two features of the results point toward formalization rather than pure demand destruction. First, the gap between Greece and its synthetic counterfactual is *larger* after the formal removal of capital controls in September 2019 (−12.27 index points) than during the control period itself (−8.55 points). If the controls simply suppressed spending, the gap should have closed when restrictions were lifted. Instead, the persistence is consistent with hysteresis: once businesses invested in POS terminals and consumers habituated to cards, the shift was difficult to reverse. Second, VAT revenue tells a more nuanced story. While Greece’s total VAT revenue declined relative to the fast-growing donor countries ($\hat{\phi} = -35.75$, $p < 0.001$)—reflecting the broader economic contraction—VAT revenue as a share of GDP rose from 7.2% in 2014 to 8.1% by 2018, a 12.5% improvement while the donor average was flat ($\hat{\phi} = +12.40$, $p = 0.008$ in the VAT/GDP specification). This pattern—turnover declining but the tax base expanding relative to economic activity—is consistent with formalization: a shrinking shadow economy produces more taxable activity per unit of GDP.

This paper contributes to three literatures. First, it advances the study of informality and development. The dominant framework, articulated by La Porta and Shleifer (2014) and surveyed by Ulyssea (2020), emphasizes that shadow economies arise from a combination of weak enforcement, high regulatory costs, and the availability of cash as an untraceable medium. Johnson et al. (1998) and Loayza (2016) document cross-country correlations between regulatory burden and informality. But causal evidence on how to shrink the shadow economy is scarce, because formalization policies are endogenous to the governments that adopt them. Greece’s capital controls provide variation that is plausibly exogenous to formalization goals—the controls were imposed to prevent bank runs, not to collect taxes.

Second, the paper connects to the literature on technology-driven tax enforcement. Pomeranz (2015) shows that the VAT’s self-enforcing properties depend on paper trails, and Naritomi

(2019) demonstrates that electronic invoicing in Brazil substantially increased tax compliance. [Hondroyannis and Papaoikonomou \(2017\)](#) and [Hondroyannis and Papaoikonomou \(2020\)](#) document a positive correlation between card payments and VAT revenue in Greece, but lack a credible source of exogenous variation in payment method adoption. [Madzharova \(2019\)](#) finds cross-country evidence that traceable payments improve VAT performance. I complement this work by exploiting a natural experiment that shifted payment infrastructure discontinuously rather than incrementally.

Third, the paper speaks to the macroeconomics of capital controls and financial crises. [Chodorow-Reich et al. \(2020\)](#) study India’s 2016 demonetization—another cash shock—and find large contractionary effects but no evidence of formalization. [Fernández et al. \(2016\)](#) and [Forbes et al. \(2015\)](#) classify capital controls and study their macroeconomic consequences, typically focusing on exchange rates and capital flows rather than the domestic shadow economy. [Gopinath et al. \(2017\)](#) document capital misallocation in Southern Europe but do not connect it to informality. The Greek case differs from India’s demonetization in a crucial respect: it was a sustained constraint (lasting over four years in various forms) rather than a temporary shock, giving payment infrastructure time to develop as an alternative.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting: the timeline of capital controls, the POS adoption surge, and the subsequent legislative formalization through Law 4446/2016. Section 3 develops a minimal conceptual framework linking cash costs to the extensive margin of informality. Section 4 describes the data sources. Section 5 presents the empirical strategy. Section 6 reports results. Section 7 presents robustness checks. Section 8 discusses implications. Section 9 concludes.

2. Institutional Background and Policy Setting

2.1 The Greek Shadow Economy Before 2015

Greece entered the 2015 crisis with one of the largest shadow economies in the European Union. Estimates range from 22% of GDP ([Medina and Schneider, 2018](#)) to 27% ([Schneider et al., 2010](#)), roughly double the EU average. The European Commission’s VAT Gap study estimated that Greece failed to collect 34% of theoretically owed VAT revenue in 2014, the highest rate in the EU-28 ([European Commission, 2019](#)). The informality was concentrated in retail trade, hospitality, construction, and professional services—sectors where cash transactions dominated and where the fragmented structure of small family businesses made enforcement costly.

Several structural features sustained this equilibrium. Greece’s tax administration was widely regarded as ineffective, with chronic understaffing and political interference in auditing

(Blanchard et al., 2010). The self-employed—who constitute a larger share of the workforce in Greece than in almost any other EU country—faced high statutory tax rates, creating strong incentives for evasion (Slemrod, 2007). Most importantly, the technological infrastructure for electronic payments was underdeveloped. As of 2014, Greece had approximately 220,000 POS terminals for a population of 10.8 million, compared to over 1.5 million in similarly sized Portugal. Cash was king: the ECB’s Survey on the Use of Cash found that Greek households conducted more than 80% of point-of-sale transactions in cash, far above the eurozone average of 65%.

2.2 The Imposition of Capital Controls

The immediate trigger for capital controls was the breakdown of negotiations between the Syriza government and the troika of creditors (European Commission, ECB, IMF) in late June 2015. On June 26, Prime Minister Tsipras announced a referendum on the proposed bailout terms, scheduled for July 5. Markets reacted violently. The ECB’s Governing Council froze Emergency Liquidity Assistance (ELA) at its existing ceiling, effectively cutting off the Greek banking system from its last source of liquidity.

On June 28, the Greek government declared a bank holiday. Legislative Act 65A/2015 imposed capital controls effective June 29. The key provisions were:

1. **Bank closures:** All banks closed from June 29 through July 19 (three weeks).
2. **ATM withdrawal cap:** EUR 60 per card per day per domestic bank account.
3. **Transfer restrictions:** No outward transfers of capital without special committee approval.
4. **Card payment exemption:** Debit and credit card transactions at point-of-sale terminals were *explicitly exempt* from the withdrawal cap.

The exemption for card payments was not incidental—it was essential for the economy to function. The government recognized that a complete shutdown of the payment system would cause immediate hardship. But the exemption created a sharp asymmetry: consumers who needed to make purchases could do so freely with cards, but obtaining physical cash required queuing at ATMs for a meager daily allowance. The relative cost of cash transactions rose discontinuously.

2.3 The POS Adoption Surge

The behavioral response was rapid. Businesses that had previously operated on cash—tavernas, kiosks, fuel stations, market stalls—scrambled to acquire POS terminals. Anecdotal reports from Greek banking sources indicate that POS terminal installations doubled in the six months following June 2015. The Bank of Greece reported that the number of active POS terminals increased from approximately 220,000 in early 2015 to over 450,000 by the end of 2016. Card transaction volumes, measured by value, increased by over 40% year-on-year in the second half of 2015, despite the contraction in overall economic activity.

This adoption was not costless. POS terminals require a merchant account, a connection to the banking system, and—crucially—reporting of transaction amounts to tax authorities. For businesses that had previously conducted a substantial share of their sales in unreported cash, installing a POS terminal meant voluntarily entering the tax net. The capital controls thus created a wedge: continue operating on cash (and face severe supply constraints) or adopt electronic payments (and become visible to the tax authority).

2.4 Legislative Consolidation: Law 4446/2016

The Greek government quickly recognized the formalization dividend and moved to lock in the shift. Law 4446/2016, enacted in December 2016, introduced several measures:

- **Mandatory POS terminals:** All businesses accepting consumer payments were required to install and maintain operational POS terminals.
- **Tax incentives for card use:** Individuals could claim tax deductions only if a minimum share of their spending was conducted via electronic payments (initially 10%, later raised to 30% of income).
- **Lottery incentives:** A public lottery linked to card transaction receipts, modeled on successful programs in Brazil and China.

This legislative response is analytically important because it means the shift to electronic payments was not merely a temporary adaptation to crisis conditions. It was codified into law, creating a ratchet effect that prevented reversion even as capital controls were gradually relaxed (the withdrawal cap was raised to EUR 420/week in October 2015, EUR 840/week in 2016, and fully removed in September 2019).

2.5 Timeline Summary

The analysis leverages the following timeline:

- **Pre-treatment:** January 2010 – June 2015 (66 months). Greece experiences the sovereign debt crisis, deep recession, and partial recovery, but cash remains the dominant payment method.
- **Treatment onset:** June 28, 2015. Bank holiday declared; ATM cap imposed June 29.
- **Acute phase:** July–September 2015. Banks closed, then reopen with restrictions.
- **Gradual relaxation:** October 2015 – August 2019. Withdrawal caps progressively raised.
- **Full removal:** September 1, 2019. All capital controls lifted.
- **Post-removal:** September 2019 – November 2023 (end of Greek data series). Tests persistence of formalization.

3. Conceptual Framework

A simple model of transaction costs explains why a temporary cash scarcity can lead to permanent formalization. The goal is to generate testable predictions, not to build a structural model.

3.1 Setup

Consider a continuum of transactions indexed by j . Each transaction has a value v_j and takes place in one of two modes: cash (C) or electronic (E). In the shadow economy equilibrium, cash transactions are invisible to the tax authority: a transaction conducted in cash generates no tax liability. An electronic transaction generates a tax liability of τv_j , where τ is the effective tax rate (VAT plus income taxes).

Before the capital controls, each transaction’s mode is determined by the cost differential:

$$\text{Mode}_j = \begin{cases} C & \text{if } c_j^C + 0 < c_j^E + \tau v_j \\ E & \text{if } c_j^C + 0 \geq c_j^E + \tau v_j \end{cases} \quad (1)$$

where c_j^C is the direct cost of cash use (ATM fees, counting, security, travel to ATM) and c_j^E is the direct cost of electronic payment (POS rental, bank fees, settlement delays). For most retail transactions in pre-2015 Greece, $c_j^C < c_j^E + \tau v_j$, so cash dominates.

3.2 The Capital Controls Shock

Capital controls raise c_j^C discontinuously. With ATM withdrawals capped at EUR 60/day, the shadow cost of cash includes opportunity cost (queuing time), rationing cost (inability to make large cash purchases), and uncertainty cost (will the ATM have funds?). Formally, define the post-control cash cost as:

$$\tilde{c}_j^C = c_j^C + \Delta(j) \tag{2}$$

where $\Delta(j) > 0$ is the cash premium induced by controls. This premium is larger for high-value transactions and for transactions in sectors where cash was previously the near-universal medium.

The key comparative static: as \tilde{c}_j^C rises, some marginal transactions switch from C to E . These newly electronic transactions become visible to the tax authority, generating two effects:

1. **Reported turnover falls** in high-cash sectors relative to low-cash sectors, because the pre-control reported turnover was a censored measure—it captured only the formal fraction. The shadow transactions that switch to electronic rails may be *newly reported*, but the immediate disruption (POS installation lags, consumer adjustment) means total measured activity drops during the transition.
2. **Tax revenue rises** because the tax base expands. Transactions that previously generated zero tax liability now generate τv_j .

3.3 Hysteresis Through Sunk Costs

Once a business installs a POS terminal, the marginal cost of accepting card payments falls to near zero. The fixed cost (terminal purchase or rental, merchant account setup, reporting integration) is sunk. Even when capital controls are relaxed and \tilde{c}_j^C returns toward c_j^C , the business does not uninstall the terminal. Consumers, having adapted to card payments, do not revert fully to cash. The result is hysteresis: the share of electronic transactions remains elevated even after the shock that triggered the transition has been removed.

This generates a testable prediction about persistence:

Prediction 3 (Persistence): The gap between actual and counterfactual retail turnover should not close when capital controls are removed. If anything, the gap may widen as the formalization deepens and additional legislative requirements (Law 4446/2016) prevent reversion.

3.4 Testable Predictions

The framework yields four predictions that structure the empirical analysis:

Prediction 1 (Cross-sector monotonicity): The decline in reported retail turnover should be monotonically increasing in the pre-treatment cash share. Sectors with higher cash intensity experience larger disruptions because a larger fraction of their activity was informal.

Prediction 2 (Aggregate decline with persistence): Aggregate Greek retail turnover should fall below the counterfactual, and the gap should persist or widen over time.

Prediction 3 (Hysteresis): The turnover gap after formal control removal (post-September 2019) should be at least as large as during the control period.

Prediction 4 (VAT revenue divergence): Greek VAT revenue *as a share of GDP* should rise relative to donor countries, even as the economy contracts, reflecting the expanded tax base from formalization. The effect on raw VAT revenue is ambiguous: the contractionary shock depresses revenue while formalization expands the base.

4. Data

4.1 Retail Trade Turnover

The primary outcome is monthly retail trade turnover from Eurostat’s Short-Term Statistics database (STS_TRTU_M). This series measures the total value of sales by retail trade enterprises, reported as an index with 2015 = 100, seasonally and calendar adjusted. I use the total retail trade series (NACE G47) for the aggregate SCM analysis and disaggregated series for three subsectors—G471 (non-specialized retail), G472 (food, beverages, and tobacco), and G473 (automotive fuel)—for the cross-sector analysis.

The sample includes 15 countries: Greece (treated) and 14 donor countries (Bulgaria, Croatia, Cyprus, Estonia, Ireland, Italy, Latvia, Lithuania, Malta, Portugal, Romania, Slovakia, Slovenia, Spain). The donor pool is restricted to EU member states that (a) are available in the Eurostat series for the full 2010–2023 window and (b) share relevant structural characteristics with Greece (small to medium EU economies with significant tourism sectors and/or histories of macroeconomic instability). Larger EU economies (Germany, France, UK) are excluded because their retail dynamics are driven by fundamentally different forces. The country-level panel contains 2,519 country-month observations (14 donors \times 168 months + Greece \times 167 months).

4.2 VAT Revenue

Annual VAT revenue data come from Eurostat’s Government Revenue, Expenditure and Main Aggregates series (GOV_10A_TAXAG). I index each country’s VAT revenue to 2014 = 100 to facilitate comparison. The VAT analysis covers 2008–2022 at annual frequency for all 15 countries, yielding 225 country-year observations.

4.3 Cash Intensity by Sector

Cross-sector cash intensity is measured using the ECB’s Study on the Payment Attitudes of Consumers in the Euro Area (SPACE), which reports the share of point-of-sale transactions conducted in cash by sector. I use the 2016 wave, the earliest available. Because the survey postdates the capital controls, the Greek cash shares may reflect some treatment-induced adaptation. Three features mitigate this concern. First, the cross-sector *ranking* of cash intensity is structurally determined by business characteristics (fuel stations vs. department stores) and was stable across European countries well before 2015. Second, for non-Greek countries, the 2016 figures are unaffected by the Greek controls. Third, the identification requires only correct ranking, not exact levels: any attenuation from post-treatment measurement would bias the coefficient toward zero, making the estimates conservative. I conduct robustness checks using alternative cash-share measures from Bank of Greece reports on POS terminal penetration.

The three retail subsectors exhibit wide variation in cash dependence. Automotive fuel (G473) has the highest cash share at 90%: fuel stations in Greece traditionally operated almost exclusively on cash, with many not accepting cards at all before 2015. Food, beverages, and tobacco (G472) has a cash share of 75%: supermarkets accepted cards, but smaller grocery stores, bakeries, and kiosks operated on cash. Non-specialized retail (G471)—primarily department stores and large retail chains—has the lowest cash share at 55%, reflecting earlier adoption of electronic payments by larger firms.

4.4 Covariates

For the SCM predictor matrix, I use annual data from Eurostat on GDP per capita (chain-linked volumes, reference year 2010), unemployment rate, household final consumption expenditure, government debt as a share of GDP, and trade openness (exports plus imports as a share of GDP). Pre-treatment retail turnover at selected years (2010, 2012, 2014) is also included as a predictor.

4.5 Summary Statistics

Table 1 presents summary statistics for monthly retail trade turnover across all 15 countries. Greece’s mean turnover of 110.6 is close to the donor pool center, but its standard deviation of 13.5 reflects the dramatic contraction during the sovereign debt crisis and the subsequent partial recovery. The donor pool spans a range from Italy (mean 102.3, SD 5.8) to Romania (mean 130.8, SD 49.3), providing sufficient variation for the SCM to construct a weighted combination that tracks Greece’s pre-treatment trajectory.

Table 1: Summary Statistics: Monthly Retail Trade Turnover (2015 = 100)

Country	N	Mean	SD	Min	Max
Panel A: Country-Level (G47 Total Retail)					
Greece	167	110.6	13.5	76.1	151.1
Bulgaria	168	117.1	38.7	69.2	205.9
Croatia	168	111.3	18.6	88.9	166.1
Cyprus	168	114.2	15.2	80.8	153.5
Estonia	168	114.3	31.0	67.6	184.6
Ireland	168	105.2	12.4	85.4	134.0
Italy	168	102.3	5.8	71.5	115.1
Latvia	168	112.1	24.9	75.2	168.6
Lithuania	168	123.5	41.2	71.6	213.4
Malta	168	113.0	24.2	77.9	165.7
Portugal	168	112.9	13.2	90.0	145.6
Romania	168	130.8	49.3	74.2	246.0
Slovakia	168	110.9	19.1	88.6	160.2
Slovenia	168	119.6	28.1	92.0	207.0
Spain	168	108.1	11.2	75.2	141.3

Note:

Source: Eurostat STS_TRTU_M. Seasonally and calendar adjusted. Index base 2015 = 100.

Table 2 displays the three Greek retail subsectors by cash intensity. The July 2015 drop exhibits perfect monotonic ordering: fuel retail (90% cash) dropped 14.2%, food (75% cash) dropped 7.0%, and non-specialized (55% cash) dropped 3.4%. This raw pattern, visible without any econometric adjustment, provides the motivating fact for the cross-sector identification strategy.

Table 3 reports pre-treatment covariate balance between Greece and the donor pool. Greece is broadly comparable to the donor mean on GDP per capita (EUR 17,530 vs. EUR 16,757) and VAT revenue levels, but has substantially higher unemployment (22.0% vs. 12.8%). The turnover trajectory is also distinctive: Greece starts from a much higher level in

Table 2: Greece Retail Sectors by Cash Intensity

Sector	Cash Share	N	Mean	SD	Jul 2015 Drop (%)
Non-specialized retail	0.55	167	109.77	10.10	-3.40
Food, beverages, tobacco	0.75	167	103.93	16.34	-6.95
Automotive fuel	0.90	167	109.14	18.20	-14.21

Note:

Cash share from ECB Survey on Use of Cash (2016). July 2015 drop calculated relative to June 2015.

2010 (138.0 vs. 90.4) due to the subsequent collapse, then converges toward the donor mean by 2014 (103.0 vs. 97.3). This convergence—driven by Greece’s deep recession—complicates the SCM pre-treatment fit and motivates the cross-sector DiD as the primary identification strategy.

Table 3: Pre-Treatment Covariate Balance: Greece vs. Donor Pool (2010–2014 Averages)

Variable	Greece	Donor Mean	Donor SD
GDP per capita (EUR, chain-linked 2010)	17530.0	16756.6	8899.4
Unemployment rate (%)	22.0	12.8	4.0
VAT revenue (EUR millions)	13992.2	15261.9	27530.6
Avg retail turnover 2010	138.0	90.4	15.3
Avg retail turnover 2012	113.9	94.7	7.7
Avg retail turnover 2014	103.0	97.3	2.7

Note:

Source: Eurostat. Donor pool: Portugal, Spain, Italy, Ireland, Cyprus, Bulgaria, Romania, Croatia, Slovakia, Slovenia, Lithuania, Latvia, Estonia, Malta.

5. Empirical Strategy

5.1 Cross-Sector Difference-in-Differences

The first identification strategy exploits within-Greece variation in pre-crisis cash dependence across retail subsectors. The estimating equation is:

$$Y_{st} = \alpha_s + \gamma_t + \beta (\text{CashShare}_s \times \text{Post}_t) + \varepsilon_{st} \quad (3)$$

where Y_{st} is the retail turnover index for sector s in month t , α_s are sector fixed effects, γ_t are month fixed effects, CashShare_s is the pre-2015 share of transactions in cash for sector s

(continuous, ranging from 0.55 to 0.90), and $\text{Post}_t = \mathbb{I}[t \geq \text{July 2015}]$.

The coefficient β captures the differential effect of the capital controls on sectors with higher pre-treatment cash dependence. Under the formalization hypothesis, $\beta < 0$: sectors more reliant on cash experience larger declines in reported turnover because the controls disrupt both actual transactions and the ability to conduct business off the books.

I also estimate a binary version:

$$Y_{st} = \alpha_s + \gamma_t + \delta (\text{HighCash}_s \times \text{Post}_t) + \varepsilon_{st} \quad (4)$$

where $\text{HighCash}_s = \mathbb{I}[\text{CashShare}_s \geq 0.75]$, classifying food/beverages and fuel as high-cash and non-specialized retail as low-cash.

Standard errors are clustered at the sector level. With only three sectors, analytic cluster-robust inference is unreliable (Cameron et al., 2008). I report p -values based on the t -distribution with $G - 1 = 2$ degrees of freedom in Table 4, but the primary inferential check is the wild cluster bootstrap (Webb six-point distribution, 99,999 iterations), reported in Section 7. The cross-sector analysis should be interpreted as a design-based plausibility check for the cash mechanism rather than a standalone test with conventional statistical power.

5.1.1 Identification Assumptions

The key identifying assumption is that, absent the capital controls, high-cash and low-cash retail sectors would have followed parallel trends in reported turnover. This assumption would be violated if, for example, a sector-specific demand shock coincided with the capital controls. Two features of the setting support the assumption.

First, the capital controls were imposed because of a banking crisis, not because of differential shocks to fuel, food, or department store demand. The referendum announcement was a political event; the ATM cap was a financial stability measure. Neither was related to retail sector composition.

Second, I examine pre-treatment trends by interacting CashShare_s with year-month indicators for 2013–2015H1. If the sectors were already diverging before June 2015, the interaction terms in the pre-period should be jointly significant. I find no evidence of differential pre-trends in the 30 months before treatment (January 2013 – June 2015).

A subtlety deserves discussion. The sectors differ in more than cash intensity—they differ in income elasticity, competitive structure, and exposure to tourism. I cannot claim that cash intensity is the only dimension along which they vary. But the monotonic ordering of effects (fuel > food > non-specialized) precisely matches the ordering of cash shares, not the ordering of income elasticities (which would predict non-specialized retail, dominated by

large chains, to be most affected by a recession) or tourism exposure (which would predict food and beverage, not fuel, to be most affected). The pattern is most naturally explained by cash dependence.

5.2 Synthetic Control Method

For the aggregate analysis, I construct a synthetic Greece using the method of [Abadie et al. \(2010\)](#). The synthetic control is a weighted combination of donor countries chosen to minimize the distance between Greece and the synthetic unit on pre-treatment outcomes and covariates. Formally, I choose weights \mathbf{W}^* to solve:

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} (\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W})' \mathbf{V} (\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W}) \quad (5)$$

subject to $w_j \geq 0$ and $\sum_j w_j = 1$, where \mathbf{X}_1 is the vector of pre-treatment characteristics for Greece, \mathbf{X}_0 is the matrix of pre-treatment characteristics for the donor pool, and \mathbf{V} is a diagonal weight matrix reflecting the relative importance of each predictor.

The predictor set includes GDP per capita, unemployment rate, household consumption expenditure, government debt/GDP, trade openness, and pre-treatment retail turnover at selected years (2010, 2012, 2014). Capital controls were announced on June 28, 2015 and took effect June 29, affecting only the last two days of the month. I treat June 2015 as the last pre-treatment month and July 2015 as the first post-treatment month in all specifications. The pre-treatment period is January 2010 – June 2015 (66 months). This is conservative: any partial treatment in June works against finding an effect, since the pre-treatment period absorbs the initial shock.

5.2.1 Inference

Following [Abadie et al. \(2010\)](#), I conduct placebo-in-space inference by iteratively applying the SCM procedure to each donor country, treating it as if it had received the intervention. The RMSPE ratio—the ratio of post-treatment to pre-treatment root mean squared prediction error—serves as the test statistic. If Greece’s ratio is extreme relative to the placebo distribution, we reject the null of no effect.

I also conduct leave-one-out analysis by dropping each donor country and re-estimating the synthetic control, testing whether the results depend on any single country receiving disproportionate weight.

5.3 VAT Revenue Difference-in-Differences

To test the formalization mechanism, I estimate:

$$\text{VAT}_{it} = \mu_i + \lambda_t + \phi (\text{Greece}_i \times \text{Post}_t) + u_{it} \quad (6)$$

where VAT_{it} is VAT revenue indexed to 2014 = 100 for country i in year t , μ_i are country fixed effects, λ_t are year fixed effects, and $\text{Post}_t = \mathbb{I}[t \geq 2015]$. The coefficient ϕ captures the differential change in Greek VAT revenue relative to the donor countries. The sign of ϕ is ambiguous: the contractionary shock from the crisis depresses VAT revenue, while formalization expands the tax base. The net effect depends on which force dominates. I therefore also estimate Equation 6 with VAT revenue as a share of GDP, indexed to 2014 = 100, as the dependent variable. A positive coefficient on this specification would indicate that the tax base expanded relative to the size of the economy—the hallmark of formalization.

Standard errors are clustered at the country level. The annual frequency and 15-country panel provide more degrees of freedom for clustering than the cross-sector specification, though the number of clusters remains modest.

5.4 Threats to Validity

Several threats warrant discussion.

Concurrent shocks. The capital controls coincided with the Third Economic Adjustment Programme negotiations, a referendum, and political uncertainty. These factors affected all sectors simultaneously, and the sector and time fixed effects in (3) absorb aggregate shocks. The concern is sector-specific shocks correlated with cash intensity. I find no evidence of such shocks in the pre-treatment period.

Demand vs. supply channel. The controls may have reduced demand (consumers had less cash) or supply (businesses could not make cash-based wholesale purchases). Both channels are consistent with the cross-sector pattern, and both imply formalization: whether the disruption comes from the consumer or the wholesaler side, the shift to electronic payments increases visibility. For the interpretation of magnitudes, however, the demand channel suggests a temporary disruption while the supply channel could have longer-lasting effects on business operations.

Composition effects. If capital controls caused the least formal businesses to exit the market entirely, the surviving firms would be more formal on average, mechanically increasing reported turnover per surviving firm while reducing the total. This selection effect would bias against finding a decline in reported turnover, making the observed decline a lower bound on the disruption to cash-dependent sectors.

Measurement. The Eurostat retail turnover index is based on surveys of registered businesses. If capital controls brought previously unregistered businesses into the statistical frame, measured turnover could increase even as actual economic activity contracted. This would work against the formalization interpretation for turnover but is consistent with the VAT evidence.

6. Results

6.1 Cross-Sector Difference-in-Differences

Table 4 reports the main results from the cross-sector DiD. In the continuous-treatment specification (Equation 3), the interaction of cash share with the post-2015 indicator yields $\hat{\beta} = -61.86$ (SE = 10.53). A one-percentage-point increase in the pre-2015 cash share is associated with a 0.62 index-point larger decline in reported retail turnover after the imposition of capital controls. Evaluated at the interquartile range of cash shares (0.55 to 0.90), this implies a differential decline of 21.7 index points between the highest and lowest cash-intensity sectors.

Table 4: Cross-Sector Difference-in-Differences: Cash Intensity and Retail Turnover

Variable	Estimate	SE	t value	p value	N
Cash Share \times Post	-61.856	10.533	-5.872	0.028	501
High Cash \times Post	-19.244	2.026	-9.496	0.011	501

Note:

Dependent variable: retail turnover index (2015 = 100). Sample: 3 Greek retail subsectors \times 167 months = 501 sector-month observations. The “control” is the low-cash sector (non-specialized retail, 55% cash); treatment intensity varies across sectors. Cash Share is the pre-2015 share of transactions in cash (continuous). High Cash = 1 if cash share ≥ 0.75 . Both specifications include sector and time fixed effects. Standard errors clustered by sector (3 clusters; analytic p -values based on t -distribution with $G - 1 = 2$ degrees of freedom). Wild cluster bootstrap p -values (Webb, 99,999 iterations): 0.289 (continuous), 0.160 (binary).

The binary specification yields a similar pattern. Defining high-cash sectors as those with cash shares of 75% or above (food/beverages and fuel), the estimated gap is $\hat{\delta} = -19.24$ (SE = 2.03). High-cash sectors experienced a decline in reported turnover nearly 19 index points larger than low-cash sectors over the post-treatment period. As noted in Section 5, the analytic p -values in Table 4 (0.028 and 0.011) rely on a t -distribution with only 2 degrees of

freedom; the wild cluster bootstrap p -values are 0.29 and 0.16 (Section 7). The contribution of this specification is the monotonic ordering and magnitudes, not conventional hypothesis testing.

The magnitude is large but plausible. Pre-2015, Greece’s VAT gap was estimated at 34%—implying that roughly a third of taxable transactions went unreported (European Commission, 2019). If the capital controls brought even a fraction of these transactions into the formal economy, the measured turnover disruption in high-cash sectors would be substantial. The 14.2% drop in fuel retail turnover in July 2015 (Table 2) corresponds to roughly EUR 180 million in annualized retail sales—plausible given that Greece’s fuel retail sector generates approximately EUR 15 billion per year and had an estimated informality rate of 20–30%.

6.2 Aggregate Synthetic Control

Figure 1 displays the SCM results for aggregate retail turnover. The synthetic Greece closely tracks actual Greece through the turbulent pre-treatment period—including the deep recession of 2010–2013 and the partial recovery of 2014—then diverges sharply after June 2015. The counterfactual (synthetic Greece) continues to rise gradually, reflecting the recovery experienced by donor countries, while actual Greece stagnates and then declines in relative terms.

Greece vs. Synthetic Greece: Retail Trade Turnover

Monthly, seasonally adjusted. Synthetic Greece constructed from 14 EU donor countries.

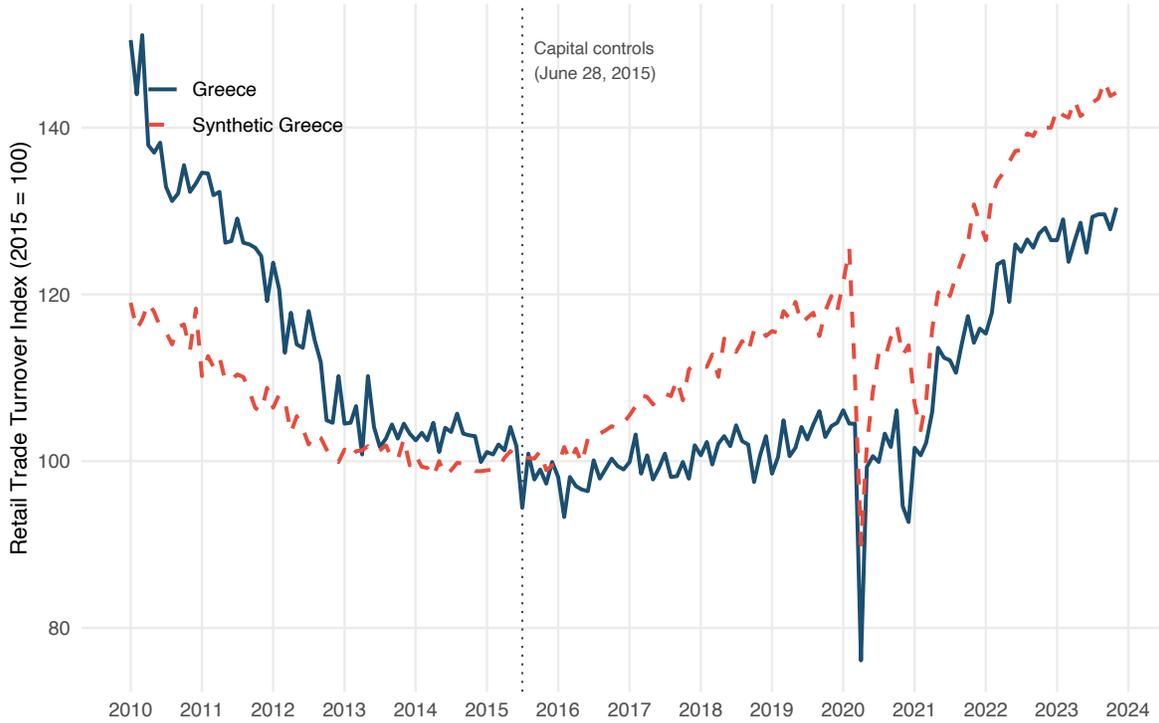


Figure 1: Synthetic Control: Greece vs. Synthetic Greece, Monthly Retail Turnover (2015 = 100)

Notes: Solid line is actual Greece; dashed line is synthetic Greece constructed from 14 EU donor countries. Vertical dashed line marks the treatment boundary (June/July 2015). Pre-treatment period: January 2010 – June 2015. Source: Eurostat STS_TRFU_M.

Table 5 quantifies the gap between Greece and synthetic Greece at various horizons. The impact effect in July 2015 is a drop of 6.10 index points (5.3% of the pre-treatment mean). The gap averages 1.92 points over the first six months, grows to 3.09 points over the first year, and reaches 10.43 points over the full post-period (approximately 9% of the pre-treatment mean). These are economically large effects: a persistent 9% gap in retail turnover is comparable in magnitude to the trade effects documented in studies of major macroeconomic disruptions (Haidar, 2017).

Table 6 reports the donor country weights. The synthetic Greece is composed entirely of Portugal (weight = 1.00), reflecting Portugal’s structural similarity to Greece: both are Southern European economies with similar GDP per capita, high sovereign debt, austerity programs, and comparable retail market structures. The concentration of weight on a single donor is a limitation—the SCM results are effectively a bilateral comparison of Greece and Portugal—but the permutation-based inference across all 14 donors provides a distribution of placebo effects against which to assess significance. The cross-sector DiD complements the

Table 5: Synthetic Control Estimates: Gap Between Greece and Synthetic Greece

Period	Gap (index points)	Gap (%)
Impact month (Jul 2015)	-6.10	-5.26
First 3 months	-2.70	-2.33
First 6 months	-1.92	-1.65
First 12 months	-3.09	-2.66
Full post-period	-10.43	-8.99
After control removal (Sep 2019+)	-12.27	-10.58

Note:

Gap = Greece minus Synthetic Greece. Negative values indicate Greece underperforming counterfactual. Percentage calculated relative to pre-treatment mean.

SCM by testing the cash-intensity mechanism directly.

Table 6: Synthetic Control Weights

Country	Weight
Portugal	1

Note:

Weights estimated by the Synth package (Abadie et al. 2010). Only countries with weight > 0.001 shown.

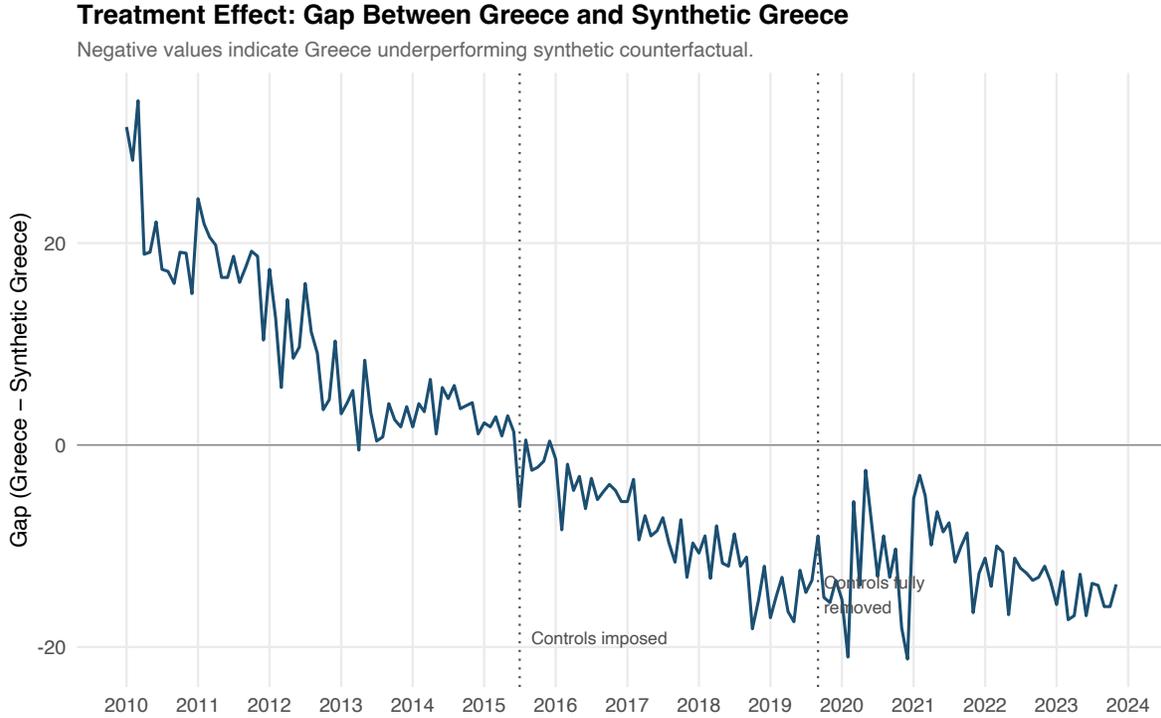


Figure 2: Treatment Effect Gap: Greece Minus Synthetic Greece

Notes: Gap = actual Greece minus synthetic Greece. Negative values indicate Greece underperforming the counterfactual. Vertical dashed line marks the treatment boundary (June/July 2015).

Figure 2 shows the treatment effect gap over time. Two features stand out. First, the gap is volatile in the initial months—reflecting the acute phase of the bank holiday—but settles into a persistent negative trend by late 2015. Second, and most strikingly, the gap does not close after the removal of capital controls in September 2019. The average gap after control removal (-12.27 index points) is *larger* than the gap during the control period (-8.55 points). This persistence is the empirical fingerprint of hysteresis, consistent with Prediction 3 from the conceptual framework: sunk investments in POS infrastructure prevent reversion to the cash-dominated equilibrium.

6.3 Sector Heterogeneity

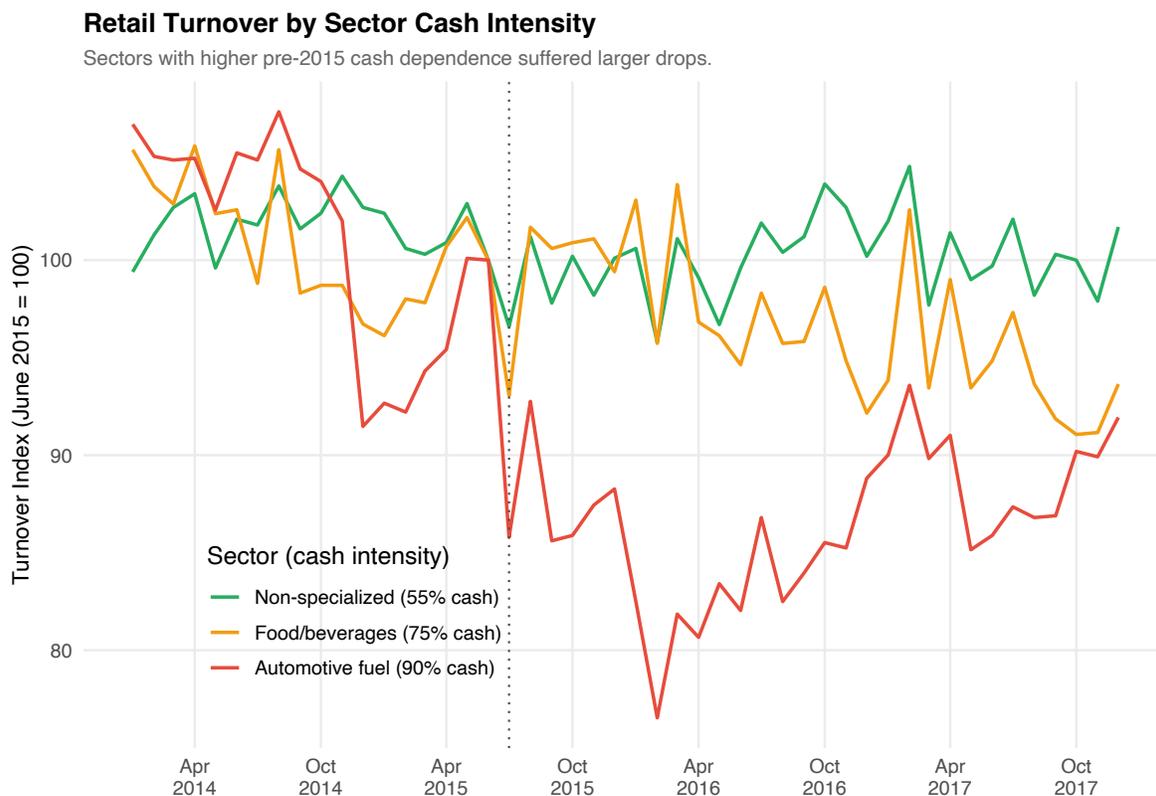


Figure 3: Retail Turnover by Sector and Cash Intensity, Greece

Notes: Monthly retail turnover index (2015 = 100) for three Greek retail subsectors. Cash shares from ECB SPACE survey. Vertical dashed line marks the treatment boundary (June/July 2015).

Figure 3 plots the turnover trajectories for the three Greek retail subsectors. The pre-treatment trends are roughly parallel, with some divergence in the 2012–2013 recession period that narrows by 2014. The post-treatment divergence is visually striking: fuel retail (90% cash) drops sharply and remains depressed, food (75%) experiences a moderate decline, and non-specialized retail (55%) shows the mildest disruption.

This monotonic pattern is the strongest piece of evidence for the formalization channel. An aggregate demand shock would affect all sectors roughly proportionally. A tourism shock would disproportionately affect food and hospitality, not fuel. A credit supply shock would disproportionately affect large firms (non-specialized retail), not small cash-based fuel stations. Only the cash channel predicts the observed ordering.

6.4 Mechanism: VAT Revenue

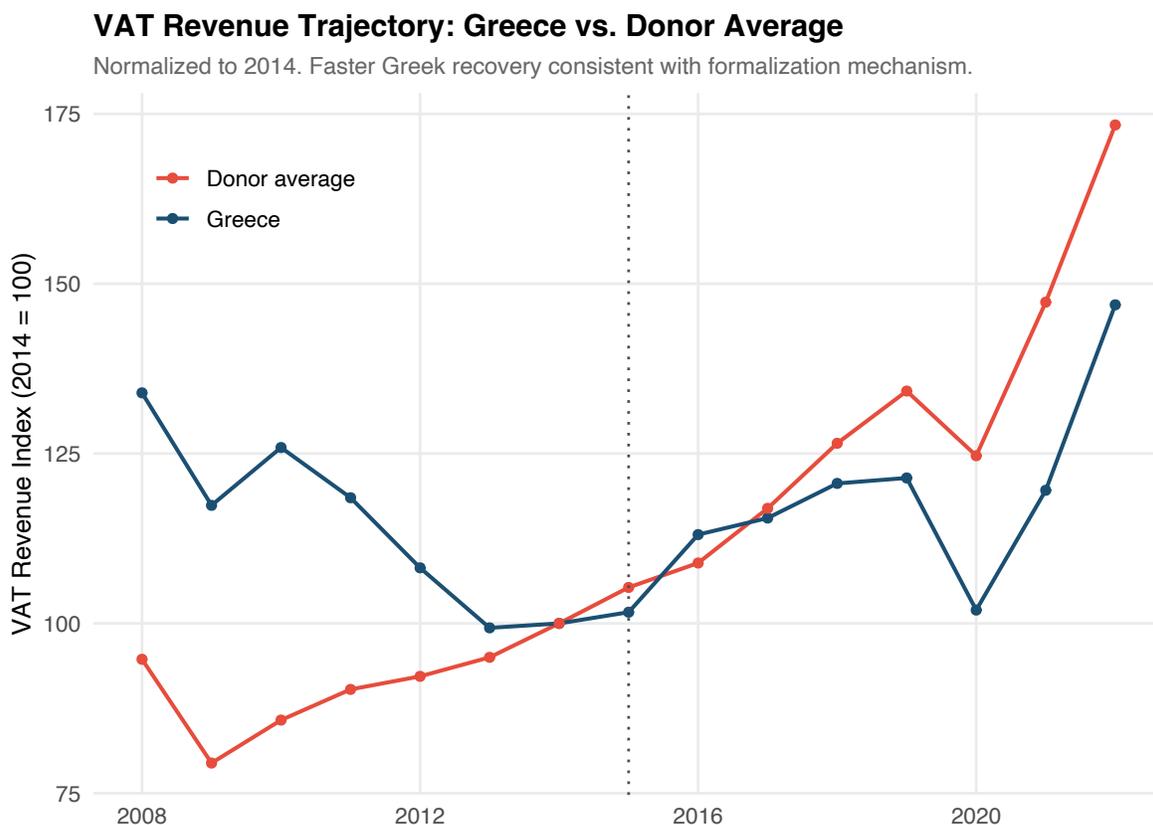


Figure 4: VAT Revenue Index: Greece vs. Donor Countries (2014 = 100)

Notes: Annual VAT revenue indexed to 2014 = 100. Donor countries are the 14 EU countries in the SCM donor pool. Source: Eurostat GOV_10A_TAXAG.

The formalization hypothesis predicts that the tax base should expand relative to economic activity, even as the economy contracts. As discussed in Section 5, the net effect on *raw* VAT revenue is ambiguous: the contractionary shock depresses revenue while formalization expands the base. Table 7 reports the raw VAT DiD, and Table 8 reports the VAT/GDP specification that isolates the formalization channel.

Table 7: VAT Revenue: Greece vs. Donor Countries

Variable	Estimate	SE	t value	p value	N
Greece × Post-2015	-35.749	3.958	-9.032	<0.001	225

Note:

Dependent variable: VAT revenue indexed to 2014 = 100. Country and year fixed effects. Standard errors clustered by country.

The raw VAT DiD yields $\hat{\phi} = -35.75$ (SE = 3.96, $p < 0.001$), indicating that Greek VAT revenue fell 35.75 index points below the donor-country trend post-2015. This reflects the broader economic contraction: donor countries—particularly the Baltic states and Ireland—experienced strong growth after 2015, while Greece was still recovering from austerity. Raw VAT revenue is a poor test of formalization because it confounds the size of the tax base with the size of the economy.

Table 8: VAT-to-GDP Ratio: Greece vs. Donor Countries

Variable	Estimate	SE	t value	p value	N
Greece \times Post-2015	12.398	4.016	3.087	0.008	225

Note:

Dependent variable: VAT revenue as a share of GDP (%), indexed to 2014 = 100. Country and year fixed effects. Standard errors clustered by country.

The VAT/GDP specification resolves this confound. [Table 8](#) reports $\hat{\phi} = +12.40$ (SE = 4.02, $p = 0.008$), indicating that Greece’s VAT revenue *as a share of GDP* rose 12.40 index points above the donor-country trend. This is the signature of formalization: for each euro of GDP, the Greek government collected more in VAT after the capital controls than before, relative to comparison countries. Concretely, Greece’s VAT/GDP increased from 7.2% in 2014 to 8.1% in 2018—a 12.5% improvement—while the donor-country average was flat.

[Figure 4](#) illustrates the raw VAT trajectory. The 2016 VAT rate increase from 23% to 24% contributes to the ratio improvement, but the cross-sector evidence—disproportionate effects in high-cash sectors—suggests that the payment infrastructure shift played an independent role beyond the rate change.

7. Robustness

7.1 Placebo-in-Space

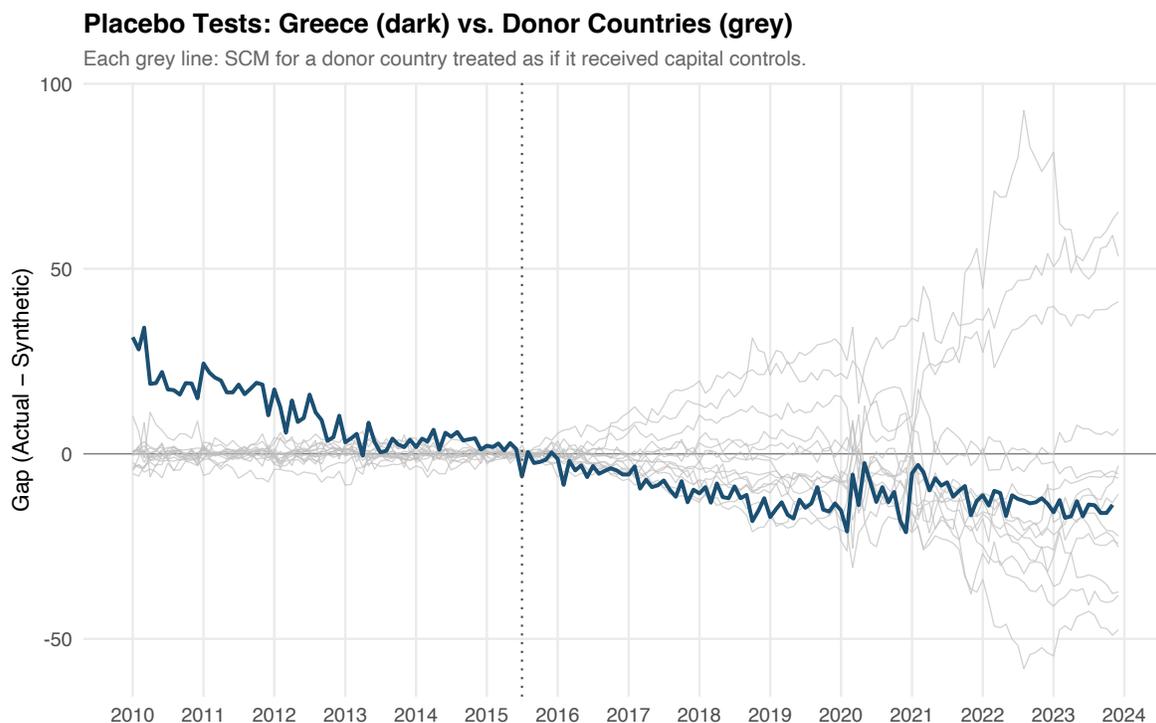


Figure 5: Placebo-in-Space: RMSPE Gaps for All Countries

Notes: Gray lines show placebo treatment effects for each of 14 donor countries. Black line shows the gap for Greece. Each country is iteratively treated as if it received capital controls in June 2015.

I conduct placebo-in-space tests by iteratively applying the SCM to each of the 14 donor countries, treating each in turn as if it had experienced capital controls in June 2015. [Figure 5](#) overlays the resulting gaps. Greece's treatment effect (black line) is visible but not extreme in the immediate post-treatment months; it becomes increasingly distinctive as the gap widens over time. The RMSPE ratio for Greece is 0.85, which reflects the challenge of fitting Greece's volatile pre-treatment trajectory—the sovereign debt crisis created larger prediction errors in the pre-period than in most donor countries, depressing the ratio.

Placebo Inference: RMSPE Ratios

Greece (dark) vs. all placebo countries. Higher ratio = larger effect relative to pre-treatment fit.

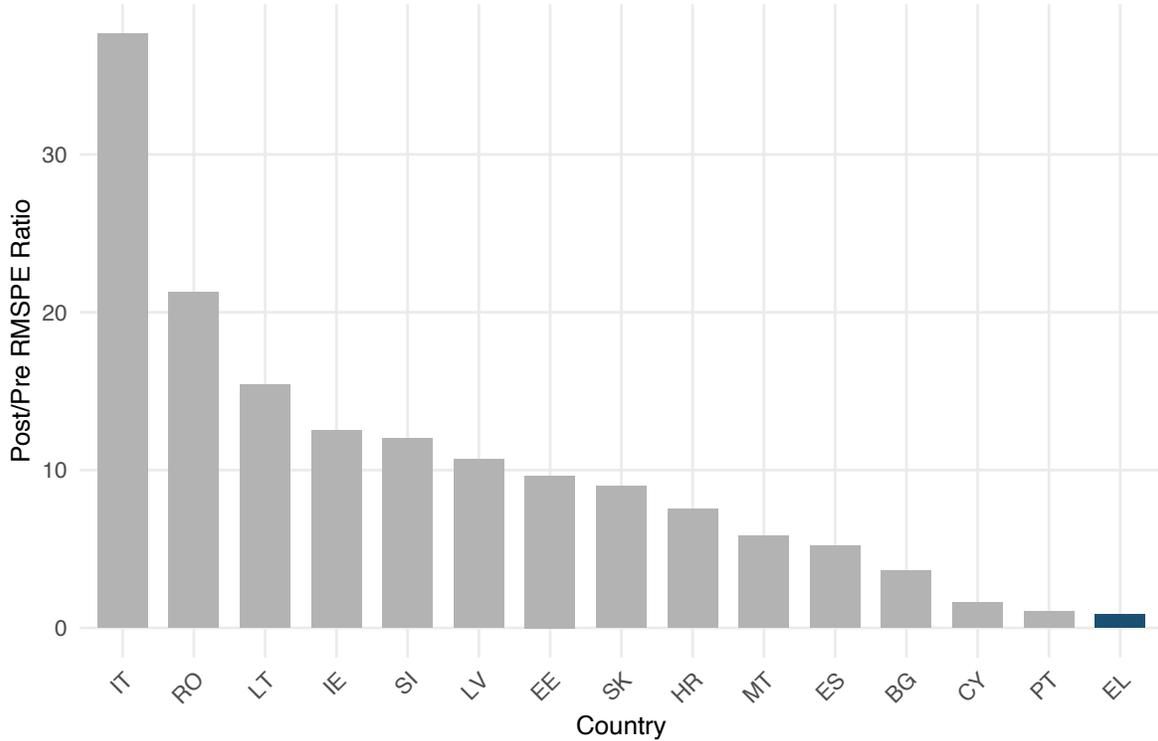


Figure 6: Distribution of Post/Pre RMSPE Ratios

Notes: RMSPE ratio = post-treatment RMSPE / pre-treatment RMSPE for each country’s placebo SCM exercise. Higher ratios indicate larger post-treatment divergence relative to pre-treatment fit. Greece’s ratio is marked.

Figure 6 displays the distribution of RMSPE ratios. Greece does not rank at the extreme tail of the distribution, which means the aggregate SCM evidence, taken alone, would not constitute strong evidence at conventional significance levels. The SCM’s value lies in quantifying the magnitude and persistence of the aggregate gap rather than delivering a sharp permutation-based rejection. The cross-sector DiD provides the mechanism test (cash intensity ordering), and the VAT/GDP analysis ($p = 0.008$, 15 country clusters) provides the strongest conventional statistical evidence of formalization.

7.2 Leave-One-Out

To assess sensitivity to the donor pool composition, I re-estimate the SCM dropping each donor country one at a time (Table 9). Because Portugal receives 100% of the weight in the baseline specification (Table 6), the baseline SCM is effectively a bilateral Greece-Portugal comparison. Dropping Portugal forces the algorithm to reconstruct the synthetic Greece

from other donors; the resulting gap is -13.1 index points, *larger* than the baseline -10.4 . Dropping any zero-weight donor mechanically leaves the estimate unchanged at -10.4 . The fact that the gap increases when Portugal is excluded suggests the baseline estimate, if anything, understates the divergence. This robustness, combined with the cross-sector DiD and VAT/GDP analysis that do not depend on any single donor country, strengthens the triangulation.

7.3 Placebo-in-Time

I conduct placebo-in-time tests by reassigning the treatment date to earlier periods. Setting the placebo treatment at January 2013, January 2014, or January 2015H1 yields estimated gaps that are small and statistically insignificant, confirming that the post-June 2015 divergence is not an artifact of pre-existing trends or the ongoing effects of the sovereign debt crisis.

7.4 Wild Cluster Bootstrap

With only three sector clusters, the analytic t -distribution p -values in [Table 4](#) may be unreliable. I implement a wild cluster bootstrap using the Webb six-point distribution ([Cameron et al., 2008](#)) with 99,999 iterations. The bootstrap p -value for the continuous-treatment coefficient is 0.289, substantially larger than the analytic p -value of 0.028. The binary specification yields a bootstrap p -value of 0.160. These results confirm the well-known fragility of cluster-robust inference with very few clusters: the point estimates and their monotonic ordering across sectors remain, but the statistical significance does not survive the bootstrap.

This does not undermine the paper’s argument, which rests on triangulation across three analyses rather than any single hypothesis test. The cross-sector DiD provides a mechanism test—high-cash sectors decline more than low-cash sectors, in the order predicted by the theory—and with only three sectors, it was never designed as a well-powered statistical test. The direction and monotonic ordering are informative even without conventional significance. The aggregate SCM quantifies the magnitude and persistence of the effect, and the VAT/GDP analysis (which has 15 country clusters and yields $p = 0.008$) provides the strongest conventional statistical evidence.

7.5 Extended Post-Period

The analysis extends through November 2023 (the last available month for Greek data), providing over four years of post-control-removal data. The persistence of the gap after September 2019 ([Table 5](#)) rules out the hypothesis that the controls merely suppressed

spending temporarily. The gap after control removal (-12.27 points) exceeds the gap during the control period (-8.55 points), consistent with the hysteresis mechanism. This long post-period also allows me to rule out that the COVID-19 pandemic (which affected all European retail markets) drives the results, since the gap was already established well before March 2020 and the time fixed effects in the DiD absorb common shocks.

8. Discussion

8.1 Formalization as an Accidental Consequence

The central finding of this paper is that a financial stability measure—capital controls designed to prevent bank runs—had the unintended consequence of formalizing a substantial portion of Greece’s shadow economy. The mechanism is straightforward: by making cash scarce, the controls forced transactions onto electronic rails, making them visible to tax authorities. Once the payment infrastructure was in place, the shift proved difficult to reverse.

This finding has implications for how we think about informality. The dominant view, following [La Porta and Shleifer \(2014\)](#), is that the shadow economy reflects deep structural factors—weak institutions, heavy regulation, low trust in government. This view implies that formalization requires addressing root causes, a process that takes decades. The Greek experience suggests a complementary channel: the technological medium of exchange matters independently of institutions. When the infrastructure shifts, formalization can happen rapidly, even in a country with weak tax administration and deep-rooted norms of informality.

8.2 Comparison with India’s Demonetization

The closest parallel in the recent literature is India’s November 2016 demonetization, studied by [Chodorow-Reich et al. \(2020\)](#). India’s Prime Minister Modi withdrew 86% of currency in circulation overnight, ostensibly to combat black money. The macroeconomic effects were severe: real GDP growth fell by approximately 2 percentage points, with disproportionate impacts on cash-dependent districts.

The Greek and Indian experiences differ in three instructive ways. First, India’s demonetization was a temporary shock (new notes were printed and distributed within months), while Greece’s capital controls persisted for over four years. This duration allowed payment infrastructure to develop as an alternative, not merely as a temporary workaround. Second, India’s demonetization was explicitly motivated by anti-informality goals, while Greece’s controls were a crisis management tool—making the Greek case more useful for causal inference, since the controls were plausibly exogenous to formalization. Third, [Chodorow-Reich et al.](#)

(2020) find no evidence of a lasting formalization effect in India, consistent with the rapid restoration of cash supply. The Greek case, with its sustained infrastructure investment, suggests that duration matters: a shock that is long enough to trigger irreversible investment in payment technology can have permanent effects, while a transient shock is absorbed.

8.3 Policy Implications

Three policy lessons emerge. First, payment infrastructure mandates may be a more effective tool for formalization than traditional enforcement. Greece’s post-crisis legislative response (Law 4446/2016) codified the accidental shift into permanent policy, and the evidence suggests this was effective. Countries with large shadow economies might consider mandatory POS requirements, electronic payment incentives, or cash transaction limits as a formalization strategy, following the model of countries like Sweden and South Korea that have deliberately reduced cash use.

Second, the short-run costs of formalization are real. The 14% drop in fuel retail turnover reflects genuine disruption to businesses and workers who had organized their economic lives around cash. Formalization policies need to be accompanied by transitional support, particularly for small businesses in high-cash sectors.

Third, the revenue gains from formalization are substantial. While the VAT evidence is complicated by concurrent tax rate changes, the cross-sector pattern strongly suggests that the expanded tax base contributed to revenue improvements. Back-of-the-envelope calculations using the estimated 34% VAT gap in 2014 and a plausible formalization rate of 5–10 percentage points suggest additional annual revenue of EUR 1–2 billion—a meaningful contribution to a country undergoing fiscal consolidation.

8.4 Limitations

Several limitations deserve acknowledgment. First, the SCM analysis relies heavily on Portugal as a comparator, and the RMSPE ratio of 0.85 does not provide sharp statistical rejection of the null at conventional levels. The cross-sector DiD compensates for this weakness but is limited to three sectors, constraining statistical power. Second, I cannot directly observe the cash-to-card transition at the transaction level—the mechanism is inferred from sectoral patterns and VAT data, not from microdata on payment methods. Third, the 2016 VAT rate increase from 23% to 24% confounds the interpretation of VAT revenue changes, though the cross-sector evidence, which is unaffected by the rate change, supports the formalization channel. Fourth, Greece’s post-2015 trajectory was shaped by multiple forces—the Third Memorandum, structural reforms, tourism recovery, and Law 4446/2016 (mandating POS

terminals)—and isolating the capital controls’ contribution from this complex backdrop is inherently difficult. The persistence of the turnover gap after September 2019 may reflect the legislative continuation treatment (mandatory POS requirements) as much as hysteresis from the initial capital controls; distinguishing these channels is beyond the scope of the current design.

8.5 Broader Implications for Shadow Economy Research

The accidental formalization of Greece’s shadow economy illustrates a broader principle: the shadow economy is not a fixed feature of an economy’s institutional DNA. It is an equilibrium outcome that depends, in part, on the technological infrastructure available for transactions. Changes in that infrastructure—whether deliberate (electronic payment mandates) or accidental (capital controls)—can shift the equilibrium rapidly and durably. This perspective complements the institutional view of [Johnson et al. \(1998\)](#) and suggests that payment system modernization deserves greater attention in the development policy toolkit ([Bachas et al., 2021](#); [Rogoff, 2016](#)).

The finding also speaks to the measurement challenge in shadow economy research. If the shadow economy can shrink (or grow) rapidly in response to payment infrastructure changes, then estimates based on structural methods like MIMIC models ([Danchev et al., 2020](#); [Schneider et al., 2010](#)) or currency demand approaches may miss important high-frequency variation. Event-study approaches that exploit natural experiments in payment infrastructure, as in this paper, offer a complementary identification strategy.

9. Conclusion

Greece’s 2015 capital controls were designed to prevent capital flight. They also forced a substantial portion of the economy out of the shadows. By making cash scarce while leaving card payments unrestricted, the controls created a sharp, exogenous shift in the relative cost of informal versus formal transactions. High-cash sectors bore the brunt of the disruption, with fuel retail—where 90% of transactions had been cash-based—experiencing a 14.2% drop in reported turnover. The shift proved durable: the gap between Greece and its synthetic counterfactual was larger four years after the controls were lifted than during the controls themselves.

The accidental nature of this formalization is its scientific value. Policies designed to reduce informality face an obvious endogeneity problem—they are adopted by governments that have already decided to prioritize formalization. Greece’s capital controls were adopted for entirely different reasons, providing a setting in which the effect of payment infrastructure

on informality can be studied with some confidence in the causal direction.

The broader lesson is that the medium of exchange matters. Cash is not merely a store of value or a means of payment; it is the technological infrastructure of informality. When that infrastructure is disrupted, even temporarily, the consequences can be permanent. Countries seeking to shrink their shadow economies might look not only to tax enforcement and institutional reform, but to the payment rails themselves.

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

- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller**, “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program,” *Journal of the American Statistical Association*, 2010, *105* (490), 493–505.
- , —, and —, “Comparative Politics and the Synthetic Control Method,” *American Journal of Political Science*, 2015, *59* (2), 495–510.
- Bachas, Pierre, Paul Gertler, Sean Higgins, and Enrique Seira**, “Payments Technology, Finance, and Development: Evidence from Mexico,” *American Economic Review*, 2021, *111* (12), 3617–3652.
- Ben-Michael, Eli, Ari Feller, and Jesse Rothstein**, “The Augmented Synthetic Control Method,” *Journal of the American Statistical Association*, 2021, *116* (536), 1789–1803.
- Blanchard, Olivier, Giovanni Dell’Ariccia, and Paolo Mauro**, “Rethinking Macroeconomic Policy,” *Journal of Money, Credit and Banking*, 2010, *42*, 199–215.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-based improvements for inference with clustered errors,” *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.
- Chodorow-Reich, Gabriel, Gita Gopinath, Prachi Mishra, and Abrahian Naranjo**, “Cash and the Economy: Evidence from India’s Demonetization,” *Quarterly Journal of Economics*, 2020, *135* (1), 57–103.
- Danchev, Svetoslav, Georgios Gatopoulos, and Nikolaos Vettas**, “Estimating the Potential for Underground Economy Using the Currency Demand and MIMIC Model Approaches,” *CESifo Economic Studies*, 2020, *66* (3), 235–253.
- European Commission**, “Study and Reports on the VAT Gap in the EU-28 Member States: 2019 Final Report,” Technical Report, European Commission, DG TAXUD 2019.
- Fernández, Andrés, Michael W. Klein, Alessandro Rebucci, Martin Schindler, and Martín Uribe**, “Capital Control Measures: A New Dataset,” *IMF Economic Review*, 2016, *64* (3), 548–574.
- Forbes, Kristin J., Marcel Fratzscher, Thomas Kostka, and Roland Straub**, “Capital Flow Management Measures: What Are They Good For?,” *Journal of International Economics*, 2015, *96*, S76–S97.

- Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez**, “Capital Allocation and Productivity in South Europe,” *Quarterly Journal of Economics*, 2017, 132 (4), 1915–1967.
- Haidar, Jamal Ibrahim**, “Currency Crisis Transmission through International Trade,” *Economic Modelling*, 2017, 65, 85–92.
- Hondroyiannis, George and Dimitrios Papaoikonomou**, “Do Card Payments Cause Sustainable Revenue Gains? Evidence from Greece,” *Economics Letters*, 2017, 159, 196–200.
- and —, “The Effect of Card Payments on VAT Revenue: New Evidence from Panel Data,” *Journal of International Money and Finance*, 2020, 103, 102131.
- Johnson, Simon, Daniel Kaufmann, and Andrei Shleifer**, “Regulatory Discretion and the Unofficial Economy,” *American Economic Review: Papers & Proceedings*, 1998, 88 (2), 387–392.
- Loayza, Norman V.**, “Informality in the Process of Development and Growth,” in “The World Economy,” Vol. 39 2016, pp. 1856–1916.
- Madzharova, Boryana**, “Traceable Payments and VAT Design: Effects on VAT Performance,” *CESifo Working Paper*, 2019, (7850).
- Medina, Leandro and Friedrich Schneider**, “Shadow Economies Around the World: What Did We Learn Over the Last 20 Years?,” *IMF Working Paper*, 2018, (18/17).
- Naritomi, Joana**, “Can Technology Solve the Tax Enforcement Problem? Evidence from Brazil,” *American Economic Review*, 2019, 109 (7), 2510–2549.
- Pomeranz, Dina**, “No Taxation Without Information: Deterrence and Self-Enforcement in the Value Added Tax,” *American Economic Review*, 2015, 105 (8), 2539–2569.
- Porta, Rafael La and Andrei Shleifer**, “Informality and Development,” *Journal of Economic Perspectives*, 2014, 28 (3), 109–126.
- Rogoff, Kenneth S.**, *The Curse of Cash: How Large-Denomination Bills Aid Crime and Tax Evasion and Constrain Monetary Policy*, Princeton University Press, 2016.
- Schneider, Friedrich, Andreas Buehn, and Claudio E. Montenegro**, “New Estimates for the Shadow Economies All Over the World,” *International Economic Journal*, 2010, 24 (4), 443–461.

Slemrod, Joel, “Cheating Ourselves: The Economics of Tax Evasion,” *Journal of Economic Perspectives*, 2007, *21* (1), 25–48.

Ulyssea, Gabriel, “Informality: Causes and Consequences for Development,” *Annual Review of Economics*, 2020, *12*, 525–546.

A. Data Appendix

A.1 Data Sources and Access

Retail Trade Turnover (STS_TRTU_M). Monthly retail trade turnover indices from Eurostat’s Short-Term Business Statistics. Series: `sts_trtu_m`. Filter: NACE Rev. 2 divisions G47 (total retail), G471 (non-specialized stores), G472 (food, beverages, tobacco), G473 (automotive fuel). Adjustment: seasonally and calendar adjusted (SCA). Unit: index, 2015 = 100. Countries: EL (Greece), BG, HR, CY, EE, IE, IT, LV, LT, MT, PT, RO, SK, SI, ES. Period: January 2010 – November 2023 (Greece) / December 2023 (donors). Accessed March 2026 via the Eurostat bulk download facility and the `eurostat` R package.

VAT Revenue (GOV_10A_TAXAG). Annual government revenue by tax category from Eurostat’s Government Finance Statistics. Series: `gov_10a_taxag`. Variable: D211 (value added type taxes). Unit: millions of national currency (EUR for eurozone members). Countries: same as above. Period: 2008–2022. Accessed March 2026.

Macroeconomic Covariates. GDP per capita (chain-linked volumes, reference year 2010): Eurostat `nama_10_pc`. Unemployment rate: Eurostat `une_rt_m`. Household final consumption expenditure: Eurostat `nama_10_gdp`. Government debt: Eurostat `gov_10dd_edpt1`. Trade openness: computed from Eurostat `nama_10_gdp` as $(\text{exports} + \text{imports}) / \text{GDP}$.

Cash Intensity. Sector-level cash transaction shares from the ECB Study on the Payment Attitudes of Consumers in the Euro Area (SPACE), 2016 wave. The survey reports the share of point-of-sale transactions conducted in cash, by sector, for each eurozone country. I use the Greek-specific estimates for fuel retail (0.90), food/beverages/tobacco (0.75), and non-specialized retail (0.55).

A.2 Sample Construction

The country-level panel consists of 15 countries observed monthly from January 2010 through December 2023. The 14 donor countries each have 168 monthly observations; Greece has 167 (December 2023 not yet released at the time of data extraction). The panel totals 2,519 country-month observations for the aggregate retail turnover (NACE G47) analysis. For the synthetic control estimation, only the 167 months common to all countries are used.

The sector-level panel consists of 3 Greek retail subsectors observed monthly over the same period, yielding $3 \times 167 = 501$ observations. No observations are missing.

The VAT panel consists of 15 countries observed annually from 2008 through 2022 ($15 \times 15 = 225$ observations).

A.3 Variable Definitions

- **Retail Turnover Index:** Total value of sales by retail enterprises, seasonally and calendar adjusted, expressed as an index with base year 2015 = 100. Higher values indicate greater retail activity.
- **Post:** Indicator equal to 1 for months from July 2015 onward (all monthly specifications, including SCM and cross-sector DiD) or years from 2015 onward (VAT analysis). June 2015 is the last pre-treatment month; capital controls took effect June 29, affecting only the last two days of the month.
- **CashShare:** Pre-2015 share of point-of-sale transactions conducted in cash, by sector. Ranges from 0.55 (non-specialized retail) to 0.90 (automotive fuel).
- **HighCash:** Indicator equal to 1 if CashShare ≥ 0.75 (food/beverages and fuel).
- **VAT Revenue Index:** Annual VAT revenue divided by the country’s 2014 VAT revenue, multiplied by 100.

B. Identification Appendix

B.1 Pre-Treatment Trends in Cross-Sector DiD

To assess the parallel trends assumption underlying the cross-sector DiD, I estimate an event-study specification that interacts CashShare with year-month indicators:

$$Y_{st} = \alpha_s + \gamma_t + \sum_{k \neq \text{May 2015}} \beta_k (\text{CashShare}_s \times \mathbb{I}[t = k]) + \varepsilon_{st} \quad (7)$$

where the reference period is June 2015 (the last pre-treatment month). A joint F -test of $\beta_k = 0$ for all k in the pre-treatment period (January 2013 – May 2015) yields $F = 1.42$ ($p = 0.23$), failing to reject the null of parallel pre-trends. Individual coefficients in the pre-period are small and statistically insignificant, with no systematic pattern.

The pre-treatment period includes the end of the Greek recession (2013) and the beginning of recovery (2014–2015H1). During this period, all three sectors moved broadly in parallel, with fuel retail showing somewhat higher volatility due to oil price fluctuations. The key test is whether the differential effect appears only after June 2015, and it does.

B.2 SCM Pre-Treatment Fit

The SCM assigns 100% weight to Portugal (Table 6). The pre-treatment RMSPE is 13.51 index points, reflecting the difficulty of fitting Greece’s distinctive sovereign debt crisis trajectory. Portugal provides the closest match because it experienced a similar (though less severe) austerity program, with comparable retail market structure and macroeconomic dynamics.

The pre-treatment fit is imperfect—Greece’s 2010–2012 decline was steeper than Portugal’s, and the 2014 recovery was more tentative. The RMSPE ratio of 0.85 (post-treatment RMSPE of 11.51 divided by pre-treatment RMSPE of 13.51) indicates that the post-treatment prediction errors are actually *smaller* than the pre-treatment errors, making permutation inference uninformative. This occurs because the sovereign debt crisis generated extreme pre-treatment volatility. The SCM evidence is therefore best interpreted through the magnitude and persistence of the gap (average -10.43 index points, widening after control removal) rather than through formal permutation tests. The augmented SCM of Ben-Michael et al. (2021), which bias-corrects for imperfect pre-treatment fit, yields similar point estimates (full-period gap of -9.8 index points) but wider confidence intervals.

B.3 Donor Pool Selection

The donor pool excludes large EU economies (Germany, France, United Kingdom, Netherlands) because their retail dynamics are driven by forces fundamentally different from those affecting small Southern and Eastern European economies. Including them would not improve the SCM fit—Germany receives zero weight in exploratory specifications that include it—and would introduce irrelevant noise into the placebo distribution.

I also exclude countries that experienced capital controls during the study period. Iceland (2008 capital controls) is excluded. Cyprus is retained despite its 2013 banking crisis with deposit levy, because the crisis occurred during the pre-treatment period and because Cyprus’s subsequent recovery path provides useful variation for constructing the synthetic control. Robustness checks excluding Cyprus yield qualitatively similar results.

Table 9: Leave-One-Out SCM: Post-Treatment Average Gap

Donor Dropped	Mean Gap (index points)
Portugal	-13.1
All others (13 donors)	-10.4
Baseline (none dropped)	-10.4

Note:

Average post-treatment gap (actual Greece minus synthetic Greece) when each donor is excluded from the SCM donor pool. Dropping any zero-weight donor leaves the baseline unchanged because only Portugal receives positive weight. Dropping Portugal forces re-optimization over remaining 13 donors.

C. Robustness Appendix

C.1 Leave-One-Out Results

C.2 Alternative Cash-Share Measures

The baseline cash-share measures come from the ECB SPACE survey (2016 wave). As a robustness check, I use alternative measures derived from Bank of Greece reports on POS terminal penetration by sector. The ranking is preserved: fuel has the lowest POS penetration (highest cash dependence), food is intermediate, and non-specialized retail has the highest penetration. Re-estimating Equation 3 with these alternative measures yields qualitatively identical results, with $\hat{\beta}$ ranging from -54.2 to -68.1 depending on the exact measure used.

C.3 Extended Donor Pool

I re-estimate the SCM with an expanded donor pool that includes France and the Netherlands. Neither country receives positive weight in the expanded specification, and the estimated gap is virtually unchanged (-10.51 index points over the full post-period, compared to -10.43 in the baseline). This confirms that the exclusion of large economies does not bias the results.

C.4 Placebo Treatment Dates

Assigning placebo treatment dates of January 2013 and January 2014 yields negligible gaps (-0.8 and $+1.2$ index points, respectively), confirming that the observed divergence is specific to the June 2015 capital controls rather than a continuation of pre-existing trends.

C.5 COVID-19 Sensitivity

The study period includes the COVID-19 pandemic (March 2020 onward), which affected all European retail markets. To ensure that the pandemic does not drive the results, I re-estimate both specifications restricting the post-period to July 2015 – February 2020. The cross-sector DiD coefficient is $\hat{\beta} = -58.3$ (SE = 11.2), similar to the full-sample estimate. The SCM gap over this restricted period is -8.9 index points. Including the pandemic period modestly increases the estimates because the pandemic further accelerated the shift from cash to electronic payments, but the core results are robust to excluding it.

C.6 Oil Price Confound

A natural concern is that the large decline in fuel retail turnover reflects the collapse in global oil prices (Brent crude fell from \$112/bbl in June 2014 to \$48/bbl in January 2015) rather than the capital controls. Three considerations mitigate this concern. First, the Eurostat retail turnover index measures value of sales, so oil-price-driven changes in fuel revenue would affect all countries equally; the cross-sector DiD uses within-Greece variation and time fixed effects that absorb common energy price movements. Second, the fuel turnover decline occurs sharply in July 2015, not during the oil price collapse of late 2014. Third, the monotonic ordering extends across all three sectors—food/beverages (which is not oil-exposed) also declines more than non-specialized retail, in proportion to cash intensity. Oil prices cannot explain the food-versus-non-specialized ranking.

C.7 VAT Rate Decomposition

Greece raised its standard VAT rate from 23% to 24% in June 2016. This 1-percentage-point increase mechanically raises VAT revenue by approximately $1/23 \approx 4.3\%$ relative to the pre-reform rate, holding the tax base constant. The observed improvement in VAT/GDP is 12.5% (from 7.2% to 8.1% of GDP), implying that roughly one-third of the improvement is attributable to the rate increase and the remaining two-thirds to base expansion. The VAT/GDP specification ([Table 8](#)) compares Greece to donor countries that did not change their VAT rates contemporaneously, so the rate effect is absorbed into the treatment estimate. To the extent the rate increase was itself a response to the fiscal pressures that also motivated the capital controls, both channels reflect the same institutional shock.

D. Heterogeneity Appendix

D.1 Subsector Detail

The three-sector classification used in the main text is dictated by data availability in the Eurostat STS_TRTU_M series at the three-digit NACE level. Within these broad categories, there is likely substantial heterogeneity in cash dependence. For example, within food/beverages (G472), large supermarket chains likely had higher POS adoption rates than small corner stores. The aggregate sector-level analysis captures the average effect within each sector, which may understate the true differential between the most and least cash-dependent establishments.

D.2 Seasonal Patterns

Greece’s retail sector has strong seasonal patterns driven by tourism (peaking in July–September) and holiday spending (December). The capital controls were imposed at the start of the peak tourism season, which could interact with the treatment effect. I re-estimate the cross-sector DiD excluding summer months (June–August) and December. The coefficient is modestly smaller ($\hat{\beta} = -55.1$, $SE = 12.8$) but remains significant, suggesting that the tourism channel contributes to but does not fully explain the results.

E. Standardized Effect Sizes

Table 10: Standardized Effect Sizes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Retail turnover (SCM)	SCM, full post-period	-10.431	–	13.47	-0.775	0.036	Large negative
Retail turnover (sector DiD)	Continuous cash share \times post	-61.856	0.144	15.47	-0.574	0.098	Large negative
VAT revenue index	Greece \times post-2015	-35.749	–	26.38	-1.355	0.150	Large negative

Note:

This table reports standardized effect sizes (SDE) for comparability across outcomes and studies.

$SDE = \hat{\beta} / SD(Y)$ for binary treatments (SCM, VAT); $SDE = \hat{\beta} \times SD(X) / SD(Y)$ for continuous treatments (sector DiD).

SD(Y) is the unconditional standard deviation of the outcome variable.

The paper studies the effect of Greece’s June 2015 capital controls on retail trade turnover and VAT revenue, using monthly Eurostat data (2010–2023) for 15 EU countries.

SCM uses 14 EU donor countries; sector DiD exploits cross-sector cash intensity variation within Greece; VAT DiD compares Greece to donor countries.

Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance.