Louvain School of Management

Sovereign financial stress and the preservation of tangible cultural heritage in Europe

Evidence from public cultural expenditure in eleven eurozone countries between 2006 and 2024

Author: Charles Maes
Supervisor: Professor Bertrand Candelon
Academic year 2025-2026
Dissertation for the Master of Business Engineering
Master subject and focus: International Finance
Daytime schedule

Declaration

During the preparation of this master's thesis, the author used Claude (Anthropic AI) and DeepL for the following purposes:

1. Claude (Anthropic AI): The AI tool was used to clarify economic and methodological concepts, to improve academic phrasing when needed, and to provide support with the code writing.
2. DeepL: The AI tools was used to improve sentence structure and grammar, and enhance overall readability, without altering the original meaning or academic content.

After using Claude (Anthropic AI) and DeepL, the author diligently reviewed and edited all content produced by the tools. Full responsibility is taken for the final content presented in this thesis.

By signing this declaration, I affirm that the content of this master's thesis reflects my original work, augmented by the responsible and transparent use of AI.

Signed at: Le Roeulx, Belgium

Date: August 11, 2026

Signature:

Maes Charles

Acknowledgments

I would like to truly express my gratitude to Professor Candelon for his time, guidance and valuable throughout the writing of this master's thesis. I am also immensely grateful to my parents, family and friends for their belief in me, support and encouragement throughout my academic journey.

Table of contents

1. Introduction

The financial turmoil that struck Europe from 2008 onwards arrived in three stages: a banking crisis absorbed onto public balance sheets, a sovereign debt crisis, and fiscal consolidation on a scale that reshaped public budgets across the continent. The way this consolidation was composed mattered as much as its size. The burden fell disproportionately on discretionary items, which require an active annual decision to fund and are defended by no constituency comparable to those protecting entitlement programmes (Breunig & Busemeyer, 2012). Public cultural budgets sit squarely in that exposed category. For tangible cultural heritage, this exposure carries a consequence that sets it apart from most other objects of cultural policy. Conservation admits few productivity gains, since restoring a listed building requires period-appropriate materials and craftsmanship that cannot be replaced by cheaper methods without breaching the asset's protection status (Baumol & Bowen, 1966; Benhamou, 1996). A postponed restoration is therefore not a neutral deferral. Deterioration accumulates, remediation costs compound, and the adjustment channels open to other sectors are foreclosed by the regulatory framework itself. If financial crises translate into cultural budget cuts, the stock of tangible heritage absorbs them in a form that is partly irreversible. This thesis therefore asks:

How does systemic financial stress affect the public expenditure through which European states preserve their tangible cultural heritage?

The choice of subject originates in a longstanding interest in history and in the built fabric through which it remains legible. A monument that is maintained stays available to be studied and used, while one allowed to decay withdraws that possibility forever. Conservation funding is a decision about intergenerational transmission, taken annually and under fiscal constraint. Whether financial crises systematically disturb that transmission proved far less settled than the frequency of the claim suggests.

Studies isolating the effect of financial crises on cultural budgets share a recurring methodological feature that motivates the present design. While the determinants literature models cultural expenditure as a continuous function of income and labour-market conditions, the crisis itself is captured discretely through pre/post-2008 dummies, EU-IMF programme indicators, or categorical distinctions between episodes. Dummies register that a country was in crisis. They do not register how deeply. The cost of this is empirical as the most extensive cross-country evidence finds a pooled crisis dummy for cultural services statistically insignificant until national heterogeneity is explicitly modelled (Srakar et al., 2017), so that what a binary indicator detects depends heavily on how the comparison group is built. This thesis measures stress continuously instead, using the sovereign variant of the ECB's Composite Indicator of Systemic Stress (CISS; Holló et al., 2012; Garcia-de-Andoain & Kremer, 2018). The timing and relative severity of the CISS are set by financial markets rather than by the researcher's own periodisation of what counts as a crisis. Its sovereign focus also matches the transmission mechanism, from banking stress to sovereign stress to fiscal consolidation, through which crises reach public budgets.

The design rests on a balanced panel of eleven eurozone countries observed annually over 2006-2024, with cultural expenditure drawn from Eurostat's COFOG classification. The restriction to the euro area is deliberate. The transmission mechanism under test presupposes a shared monetary framework and the absence of exchange-rate adjustment as a shock absorber. Two-way fixed-effects regressions identify the effect of stress from the differential exposure of countries within a given year. Five hypotheses, formulated from the literature and frozen before estimation, are tested in turn: a direct effect on cultural expenditure as a share of GDP (H1) and of total expenditure (H2), moderation by institutional quality (H3), persistence through lagged stress (H4), and non-linearity (H5), each followed by a pre-committed robustness battery.

None of the five hypotheses is confirmed on the specification against which it was specified ex-ante. The pattern survives a battery of robustness variations, the one association that is detectable among the five hypotheses, between sovereign stress and the budgetary share of culture, is absorbed entirely by the macro-fiscal cycle. The pattern is compatible with a threshold mechanism confined to episodes of externally imposed constraint. It is equally compatible with a design too imprecise to detect a gradient, and with the attenuation the stress measure itself induces. Section 3.8 and Chapter 4 set out why this panel cannot separate the three.

Several constraints shaped the design, they do not merely qualify it. Harmonised data exist at the level of cultural services but not for heritage conservation separately, so the dependent variable is slightly broader than the object of interest. Moreover, no continuous cross-country panel documents private heritage funding, which made the substitution question untestable. A further hypothesis, on the composition of adjustment within cultural budgets, was set aside on power grounds as the minimum detectable effect on its coefficient exceeded any plausible economic magnitude by two orders of magnitude (Section 3.2). Greece, whose experience is most relevant to the question, is the one country lacking a country-level stress index. And with eleven clusters, asymptotic inference is unreliable, hence the exact bootstrap p-values and detectability bounds reported throughout. The stress index is itself rank-based within each country, so that it orders episodes within a national history rather than ranking severity across countries, a property that constrains how the coefficient is read.

2. Literature Review

2.1. Tangible cultural heritage as economic asset

2.1.1. Tangible cultural heritage as public good

The word heritage derives from the French héritage, a legal concept designating assets transmitted by hereditary succession. The term has since broadened considerably and no longer covers material assets alone but is defined by an object's capacity to embody aesthetic, historical, scientific, social and symbolic values recognised by society (Vecco, 2010). This thesis nonetheless retains the narrower definition set out in the 1972 UNESCO Convention Concerning the Protection of the World Cultural and Natural Heritage, covering the three categories set out in Article 1 of the Convention: monuments, groups of buildings and sites (UNESCO, 1972).

Like any capital item, cultural capital has an economic value, part of which is directly observable, since those who value heritage bid up the purchase and rental prices of such properties (Throsby, 2010). But heritage matters to communities for reasons no transaction records. It carries symbolic, commemorative and historical value, serving as a reminder of the duration of a shared experience and, in a sense, guaranteeing its continuation (Greffe, 2013). Throsby (2001) identifies aesthetic, spiritual and social value along the same lines. Such non-use values can be elicited through stated-preference methods such as contingent valuation, which separate the value placed on a heritage item's mere existence, the option of consuming its services at some undefined future date, and the bequest value attached to transmitting it to future generations (Rizzo & Throsby, 2006; Throsby, 2010).

Because these values accrue to non-owners, conservation generates benefits well beyond the party paying for it, and heritage assets are distributed along a spectrum running from pure public goods to club goods. The facade of a listed building approaches the pure case almost perfectly as it is enjoyed by any passer-by regardless of whether the interior is open. The consumption is therefore neither rival nor excludable and each individual may benefit while hoping others bear the cost (Samuelson, 1954; Benhamou, 2013). Paid access to the interior instead describes a club good, non-rival up to a capacity threshold beyond which congestion erodes that property (Buchanan, 1965). In either configuration the total benefit of a conservation investment is the sum of the marginal benefits accruing to all who value the asset, not merely to its owner, so that a market left to itself preserves too little.

The public-good characterisation does not justify public financing per se, since externalities could in principle be internalised through targeted local taxation instead of general subsidy. The merit-good argument supplies the complement (Musgrave, 1959, cited in Benhamou, 2013). Following this argument, goods such as education and culture are undervalued by individuals relative to their social benefit and are systematically under-provided if left to private initiative, independently of their public- or club-good characteristics.

2.1.2.The 'cost disease' as it applies to heritage

Some activities admit few productivity gains because the labour is itself the product. This structural tendency for costs to rise faster than productivity gains in labour-intensive cultural sectors is known as cost disease (Baumol & Bowen, 1966). The theory was illustrated through the example of a string quartet. It requires as many musicians and as much time to perform today as two centuries ago. So that as wages align with more productive sectors of the economy, relative costs rise steadily, pushing up prices and, where demand is price-elastic, threatening the very survival of these activities. Baumol and Bowen argued that public grants or private donations can postpone the resulting cost pressures and ensure the long-term sustainability of such activities within a market economy.

This account can, however, be questioned as a general explanation of rising cultural costs. Artistic production is not entirely exempt from technical progress, audience preferences and quality standards evolve over time, and cultural demand has in some sectors risen faster than costs; as culture behaves as a superior good, rising household income can itself offset cost increases without requiring public subsidy, and some of the observed cost growth may reflect managerial choices rather than economic inevitability. Using the cost disease as a blanket justification for public subsidy across the cultural sector as a whole may therefore appear overly simplistic (Greffe, 2010).

Built heritage, however, constitutes a case in which several of these counterarguments do not hold. An analogy between the cost disease and historic monument restoration can be drawn: as with many service industries, the technology of restoration cannot be improved without altering the nature of the product itself, since respecting the historical character of a building necessarily calls for period-appropriate materials and specialist craftsmanship (Benhamou, 1996). This technological constraint is reinforced by the regulatory framework identified in Section 2.2.1, removing one of the principal adjustment channels that Greffe invokes to question the cost disease. Given a constant, and a fortiori growing, social demand for cultural heritage, this stagnant productivity compels the state to maintain or increase the level of taxation required to fund it, regardless of whether the resulting expenditure is offset by any increase in the asset's value (Benhamou, 1996). As a result, the productivity gap between conventional construction and heritage restoration tends to widen progressively over time, and one of the few remaining means of containing costs, accepting a reduction in quality, is rarely compatible with the preservation mandate of heritage institutions. It is this combination of technological and regulatory rigidity that justifies treating built heritage as a case where the cost disease applies with particular force. This point helps reconcile Greffe's broader scepticism with the income-elasticity evidence presented in Section 2.2.3, following which state or private capital transfers from more productive sectors can offset the resulting cost pressure without resolving its structural origin, leaving heritage budgets persistently exposed when such transfers are themselves curtailed.

2.1.3. Economic benefits of heritage

Beyond its cultural significance, tangible cultural heritage represents a substantial source of economic activity at the European level, with 2.1% of total employment linked to it, 73% of which is directly tied to tourism and 25% to construction (ESPON, 2020). Its most visible economic dimension lies in its capacity to attract visitors. Tourism improves the cultural sector's economic performance, while culture strengthens the identity and branding of tourism destinations (Škrabić Perić et al., 2021). This effect is even stronger in regions where cultural tourism attractiveness is comparatively weaker (Kostakis & Lolos, 2024). Some policymakers seek to amplify these effects by pursuing UNESCO World Heritage Site designation, increasing visibility and generating substantial economic benefits at a cost of adverse effects such as over-tourism and site degradation have been documented (Cuccia, 2012; Adie et al., 2020).

These benefits are not without tensions. Their uneven distribution raises questions about how the funding burden should be redistributed across the sector. The problem lies in the fact that tourism revenue concentrates around a small number of high-visibility sites while the broader stock of less visited heritage still requires maintenance (Benhamou & Thesmar, 2011). More fundamentally, Cominelli and Greffe (2019) warn that framing heritage primarily as an economic asset risks subordinating its existence value to its use value, as illustrated by the drift of some European museums toward commercial programming at the cost of conservation. These dynamics reinforce rather than resolve the case for public funding and point to the need for sustainable financing arrangements, whose architecture across Europe is examined in the following section.

2.2. Models for funding cultural heritage in Europe

The funding awarded to cultural heritage goods in Europe relies on a diversity of sources which can be first divided into two categories: public and private. The distribution of sources differs from one member state to another, depending on the objective of their cultural policies and economic conditions. The following sections successively analyse those dimensions.

2.2.1. Public funding models in Europe

Governments have three kinds of tool: direct expenditure, tax incentives and regulation (Rizzo & Throsby, 2006). Direct expenditure addresses market failure through public ownership or through subsidies incentivising preservation. Tax incentives work on private owners instead, typically as inheritance-tax exemptions or deductions for restoration works (Benhamou, 2013). In France, for private owners of listed buildings, the rules vary depending on whether or not the building is open to visitors for a certain number of days each year (Benhamou, 2019). On the other hand, regulatory instruments establish listing criteria to determine which heritage items are sufficiently significant to warrant some public control over their use, and lay down standards for the ways in which heritage buildings and sites can be renovated, restored, altered or re-used (Throsby, 2010). An important distinction has to be made between hard regulation involving enforceable penalties and soft regulation based on non-enforceable directives (Rizzo & Throsby, 2006).

These instruments are deployed differently across Europe. The first model, found in France and Italy, relies on a strong centralised state, with a top-down system leaving final decisions to politicians and bureaucrats (Van der Ploeg, 2006). The second one is prevalent in federal states such as Germany, Switzerland or Belgium. It devolves cultural responsibility to regional or local authorities, guaranteeing socio-cultural independence and political autonomy (Freitag & Vatter, 2008). In the British case, funding runs through arm's-length Arts Councils, limiting political interference in funding allocation (Van der Ploeg, 2006). Finally, mixed models such as the Dutch system combine expert advisory councils with ministerial responsibility, balancing technical legitimacy and political accountability (Van der Ploeg, 2006). Beyond their structural differences, these models do not perform equally in times of financial stress. Overall, conducting cultural policy at as low a level of government as possible tends to improve allocative efficiency, as local politicians are better informed about their electorate's preferences. However, a competitive climate may justify retaining part of the subsidies at national or international level (Van der Ploeg, 2006). Whether decentralisation also confers resilience is more ambiguous given that the multiplicity of funding levels diversifies budgetary risk, but it can equally propagate fiscal shocks through peer effects when the shock is generalised (Dalle Nogare & Galizzi, 2011).

Despite being structurally necessary, public funding for heritage faces several fundamental constraints. Budgetary pressures arise from the steady increase in listed buildings combined with rising restoration costs. This creates a persistent gap between needs and allocated resources (Benhamou, 1996). Politically, decision-making processes remain susceptible to lobbying and rent-seeking, while bureaucrats, facing incomplete monitoring and liability risks, tend toward excessive conservationism. This situation leads to unjustifiable listing enlargement and overly strict requirements that hinder the modern use of heritage buildings (Van der Ploeg, 2006; Rizzo & Throsby, 2006). The principle of inalienability of public assets further rigidifies public expenditure by compelling the State to bear rising maintenance costs regardless of budgetary priorities (Benhamou, 2019). These combined constraints call for a rethink of heritage funding arrangements, notably through greater involvement of private actors and alternative funding mechanisms.

2.2.2. Private funding

Government spending on cultural services has declined in relative terms, averaging 1.2% of total government expenditure across OECD countries and generally remaining below its pre-2008 level since the crisis (OECD, 2022). Governments have correspondingly shifted from direct support toward an intermediary role, seeking to channel private investment to compensate. Non-governmental contributions matter all the more for built heritage because of its ownership structure. The Council of Europe (2003) estimates that around 65% of European heritage sites and buildings are privately held, so that a share of preservation costs falls on owners depending on the applicable subsidy rules. This private funding ranges from donations, patronage and sponsorship to partnerships, delegated management and partial or total privatisation (Seaman, 2013). Private ownership does not, however, only imply private financing. Listing obligations, restoration subsidies and tax expenditure mean that the public sector remains the marginal funder of conservation on privately held assets, and it is that public margin, not the total cost of preservation, that this thesis measures.

Direct private contributions come from individuals and companies, either through sponsorship, a contractual exchange in which the donor expects a return in image or visibility, or through patronage, in principle disinterested although tax deductibility and reputational benefits play a non-negligible role (Vincent & Wunderle, 2010). Most European states offer tax incentives to encourage giving, but their design varies widely from the substantial schemes of France and Italy to Belgian tax law, which grants no specific status to patronage and treats only sponsorship as a deductible professional expense (Vincent & Wunderle, 2010; OECD, 2022). On the other side, crowdfunding has emerged as a further complementary mechanism and its development varies markedly with institutional and cultural context (Jelinčić & Šveb, 2021; Di Pietro & Butticè, 2020).

These private channels share a vulnerability. Donors concentrate their support on large, established institutions and consensual works, to the detriment of smaller heritage sites; and because sponsorship and patronage are themselves discretionary expenses, they are among the first items cut in an economic downturn, mirroring on the private side the exposure of public cultural budgets and undermining any expectation that private funds will rise to offset public retrenchment precisely when it occurs (Husson, 2013).

Whether private funding substitutes for public support or complements it has been extensively debated, and the answer depends overwhelmingly on method. Experimental studies find substantial crowding-out, while observational analyses find an effect close to zero or slightly positive (De Wit & Bekkers, 2016). What matters for the present purpose is magnitude more than direction, and Dutch panel evidence indicates that substitution is at best partial, the rise in private donations failing to offset the public withdrawal (De Wit et al., 2016). A hypothesis built on this mechanism was considered and not retained, for reasons given in Section 3.2.

2.2.3.The quantitative determinants of public spending on culture.

In the cultural sector, public expenditure on culture generally follows the long-term growth trend of national income. Initial sectoral econometric analyses confirm this intuition and allow us to take the analysis further, saying that national income is the dominant explanatory variable for public cultural expenditure, with an elasticity significantly greater than one (Getzner, 2002). This result is only amplified by the 'cost disease' specific to the cultural sector (Baumol & Bowen, 1966). Public cultural expenditure therefore behaves like a superior good in the economic sense as it grows more than proportionally with income and, conversely, contracts more than proportionally when income falls. As shown in a study across Austrian provinces, the available tax base, a proxy for income, is the primary determinant of the level of cultural expenditure by states and regions, over the ideological preferences or discretionary choices of those in power (Getzner, 2015). The income elasticity established at national level therefore also governs sub-national budgets. Conversely, rising unemployment has a negative effect on cultural expenditure, as budgetary trade-offs favour other areas over culture when social pressure intensifies and other expenditure items take on greater political priority (Srakar et al., 2017). Thus, public cultural expenditure is pro-cyclical by construction.

Procyclicality is, in fact, the most directly observable empirical manifestation of this structural vulnerability. A fiscal policy is said to be procyclical when it amplifies more than stabilise the economic cycle. Expenditure rises during periods of growth and falls during recessions, running completely counter to the requirements of macroeconomic stabilisation (Gootjes & de Haan, 2020). Whilst budgetary plans often appear to be counter-cyclical, actual outcomes are consistently pro-cyclical in the Member States of the European Union (Gootjes & de Haan, 2020). However, among the expenditure items that are susceptible to adjustment, culture is particularly vulnerable, unlike spending on health which enjoys partial protection due to its mandatory nature or political prominence, making it politically more costly to cut (Breunig & Busemeyer, 2012).

This persistence of cuts calls for an institutional explanation, which the concept of the ratchet effect can provide if we reverse its logic. In its original formulation, the ratchet effect affirms that public expenditure "ratchets" during drops in the economy, mainly due to wars or social turmoils, and remains at a higher level after the economy stabilizes (Peacock & Wiseman, 1961). Cultural expenditure inverts this. Cuts made under budgetary constraint survive the crisis that produced them. An empirical test of the ratchet hypothesis over the long term in Spain concludes that it does not explain the trend in public expenditure following ordinary recessions (Jaén-García, 2021). This lack of spontaneous recovery can be explained in particular by the adjustments made under budgetary constraints, with fiscal stress prompting local authorities to increasingly outsource the provision of cultural services (Bertacchini & Dalle Nogare, 2014). Such reorganisations prove difficult to reverse once initiated.

Taken together, these factors, income elasticity, procyclicality and the entrenchment of cuts once made, paint a coherent picture of the underlying fragility of public cultural budgets. Their sensitivity to the economic cycle, amplified by their discretionary nature and the relative weakness of their political constituencies, exposes them disproportionately to budgetary shocks.

2.3. Financial crises, austerity and the resilience of heritage institutions

The chronic exposure of cultural budgets raises a direct question: how did this vulnerability translate concretely when Europe was hit by successive financial shocks since 2008? The answer turns out to depend less on the severity of each shock than on what each state's institutions could absorb.

2.3.1. Transmission mechanisms from financial crises to cultural budgets

The transmission of a financial crisis into cultural budget cuts follows a well-identified macroeconomic sequence. Banking crises are systematically followed, with a lag, by sovereign debt crises, as governments absorb private losses onto public balance sheets, leading to a sharp deterioration of fiscal positions (Reinhart & Rogoff, 2011). This sequence is precisely what Europe experienced from the banking crisis of 2008 and the sovereign debt crisis of 2011-2012, two phases of what is, from a fiscal standpoint, not two independent shocks but a single prolonged episode of stress. Once public deficits become unsustainable, governments are compelled to undertake fiscal consolidation, and the composition of this consolidation matters greatly for which sectors bear the burden.

Spending-based austerity plans are, on average, considerably less costly in terms of output loss than tax-based ones, which tend to trigger large and persistent recessions (Alesina et al., 2019). This asymmetry creates a strong incentive for governments to favour expenditure cuts over tax increases, with the burden of adjustment falling on discretionary items. Within that envelope, the choice of which items to sacrifice is itself a political economy problem: discretionary items require an active annual decision to fund and carry none of the acquired-rights protection that shields entitlement programmes such as pensions or unemployment benefits (Breunig & Busemeyer, 2012). As established in Section 2.2.3, cultural spending belongs to this exposed category, and two further features single it out even among discretionary items. It carries no minimum service standard, a municipality must run a civil registry and a primary school, but no legislation prescribes a comparable floor for cultural provision, and its funding is split across national, intermediate and local tiers, with no authority obliged to compensate a retrenchment decided elsewhere (Husson, 2013). Unlike the closure of a hospital ward, the closure of a museum or the postponement of a restoration project can also be framed as more easily reversible, at least in the short run, making culture a politically convenient target. This pattern is confirmed empirically by Dalle Nogare and Galizzi (2011). The authors found that Italian mayors reduce cultural expenditure by approximately 10% in election years while total municipal spending remains unaffected, suggesting that voters do not reward cultural investment at the ballot box.

This perceived reversibility is only partial, an asset left without maintenance accumulates deterioration whose remediation cost compounds, adding to the cost rigidity established in Section 2.1.2. This logic is complicated, however, by a second asymmetry that runs in the opposite direction. While spending-based consolidation is less costly in output terms, it is also more destabilising politically given that expenditure cuts are found to be a significantly stronger driver of social unrest than tax increases, even once the confounding effect of economic downturns is accounted for (Ponticelli & Voth, 2020). The cheaper option in growth terms is the riskier one politically. Policy-makers face a trade-off, not a straightforward incentive. One way to manage this trade-off is to favour cuts that are rapid and visible enough to limit prolonged exposure to discontent, and discretionary cultural spending satisfies this criterion more readily than reforms to entitlement programmes, whose political cost is typically more diffuse but also more enduring.

This transmission mechanism is not uniform across European states, however. As shown in Section 2.2.1, the degree of centralisation of cultural funding, whether concentrated in a single national authority or distributed across multiple levels of government, shapes how budgetary shocks propagate through the system.

2.3.2. Empirical evidence: financial crises and cultural expenditure in Europe

The mechanisms described above predict that cultural budgets should contract more sharply, and recover more slowly, in states most exposed to financial and sovereign stress. The empirical literature largely confirms this prediction, while revealing substantial heterogeneity in its magnitude and persistence across Europe.

The most comprehensive cross-country evidence is provided by Srakar et al. (2017), whose analysis of European public cultural budgets after 2008 finds statistically significant cuts at all levels of government once national and institutional heterogeneity is controlled for. This result is itself methodologically significant. Without controlling for cross-country heterogeneity, the crisis dummy for cultural services is not statistically significant in the pooled sample, suggesting that averaging across the twenty-nine countries in the study conceals dynamics that only become visible once country clusters are introduced. Crucially, the severity of these cuts is not uniform considering that countries placed under EU-IMF adjustment programmes, namely Greece, Ireland, Portugal and Cyprus, experienced markedly deeper contractions than unconstrained Western European countries, while Eastern European states displayed comparatively muted responses. Spain is a borderline case excluded from that group. Its 2012 assistance came through the European Stability Mechanism, was confined to bank recapitalisation, and left market access intact (Section 3.4). This pattern is consistent with the transmission mechanism outlined in Section 2.3.1. The states whose fiscal position forced them into externally imposed consolidation plans bore the brunt of the adjustment, since their access to debt financing, and therefore their room for manoeuvre to protect discretionary spending, was the most constrained.

The asymmetry between central and local governments is a recurrent finding in this literature. Central cultural budgets tend to contract earlier and more sharply than local ones, which absorb part of the shock through inertia or through substitution effects that displace the adjustment to lower tiers of government (Čopič et al., 2013). Srakar et al. (2017) refine this picture, showing that central budgets were predominantly affected at the onset of the crisis in 2009-2010, while local budgets experienced their most significant cuts only in 2011-2012. The adjustment was displaced downward in sequence and in size.

A North-South asymmetry is corroborated by sociological research focusing specifically on Southern Europe. Rubio-Arostegui and Rius-Ulldemolins (2020) show that, despite two decades of expansion and convergence with the rest of the continent prior to the crisis, cultural spending in Greece, Portugal, Spain and Italy remains systematically lower, relative to public budgets and to GDP, than in Northern and Nordic countries, a gap that the crisis widened instead of closed. The authors interpret this not merely as the mechanical outcome of harsher fiscal constraints, but as evidence of a distinct model of cultural policy. In this Southern model, cultural institutions remain comparatively more dependent on direct state funding and less buffered by diversified private philanthropy documented in Northern Europe, leaving them more exposed when the state retrenches.

A further empirical regularity concerns persistence. Most European countries that cut cultural spending during the crisis years have not subsequently returned to their pre-crisis funding levels (Almeda et al., 2018). Reviewing government cultural expenditure across OECD countries to 2022 on the broader COFOG recreation, culture and religion aggregate, Bonet (2026) shows that the pattern holds when the window is extended. The author finds that the states hit hardest after 2008 have only partially restored their budgetary position without regaining pre-crisis levels of financing. This lack of spontaneous recovery is consistent with the reversed ratchet effect already identified in cultural expenditure dynamics (Jaén-García, 2021; Section 2.2.3).

The COVID-19 shock of 2020 offers an instructive, if imperfect, point of comparison. The speed and scale of emergency support packages deployed across several European countries stood in sharp contrast to the slow, politically contested consolidation that followed 2008. The contrast is one of speed and reversibility rather than of magnitude, and it is on that dimension that the two episodes are not comparable (Betzler et al., 2020).

2.3.3. Institutional resilience: governance and legal frameworks

The heterogeneity documented above calls for an institutional explanation. The literature on cultural governance suggests that a national system's capacity to absorb financial shocks depends less on the severity of the shock itself than on the structural configuration of funding sources and decision-making authority that existed before the crisis struck.

Bonet and Donato (2011) offer the most direct account of what the 2008 crisis meant concretely for European cultural institutions. Their analysis shows that funding pressure fell most heavily on organisations that combined near-total dependence on a single public source with a centralised governance structure leaving little room for managerial adaptation. Independent organisations and those promoting experimental work proved especially fragile, while larger established institutions were partially shielded by their political visibility. Beyond the immediate distributional effects, the authors argue that the crisis is not a cyclical episode but a structural rupture as the welfare state model that underwrote cultural expansion over the previous half-century has been permanently reconfigured. Waiting for a return to pre-crisis funding levels is thus not a viable institutional strategy.

Where cultural institutions draw resources from multiple governmental tiers and revenue streams, a retrenchment decided at one level need not transmit fully to the others and the system as a whole can absorb part of the shock through internal redistribution rather than uniform contraction (Rizzo & Throsby, 2006; Van der Ploeg, 2006). Getzner (2015) shows that in Austria, central government expenditure partially offset provincial cultural budget cuts, but this compensatory mechanism presupposes a central level with sufficient fiscal room to absorb the shock, a condition that generalised crises, affecting all tiers of government simultaneously, are precisely those most likely to violate. Yet, the institutional configuration may matter more than multi-tier architecture in itself. Cultural spending by neighbouring jurisdictions behaves as strategic complements in Italian cities (Dalle Nogare & Galizzi, 2011) but as substitutes in Swedish ones (Lundberg, 2006). Where spending decisions reinforce one another, a common downward shock is likely to compound more than dissipate across tiers but where they substitute, the reverse may hold. The shocks at issue here are common ones. For a crisis reaching all tiers at once, the Italian regime is the more relevant of the two, and multi-tier architecture is read here as a channel that can propagate a generalised shock, not as a buffer against it. The mechanism Getzner documents in Austria works only because a single country was under strain while its neighbours were not. The federal government had spare fiscal room to draw on precisely because the shock was local. A shock that hits every tier of every government at once removes that spare room by construction. H3 therefore expects institutional quality to buffer stress less effectively than the Austrian case suggests, rather than testing the local/systemic contrast directly, since every shock observed here is of the systemic kind (Sections 3.2 and 3.4).

Prolonged fiscal pressure induced a reconfiguration of heritage governance. Albertelli and Mignosa (2026) situate the growing adoption of Public-Private Partnerships (PPPs) in European heritage governance within a broader shift in public administration models, from hierarchical state provision toward collaborative arrangements in which public bodies, private actors and non-profit organisations jointly share resources, responsibilities and risks. This shift was theorised by Osborne (2006) as the transition from traditional Public Administration and New Public Management toward New Public Governance. Periods of financial crisis have intensified this transition by widening the gap between the growing costs of managing cultural heritage and the public funding available to meet them, a dynamic formalised in Italy through the 2016 introduction of Special PPPs specifically designed for the cultural sector. Nevertheless, risks are unlikely to be evenly shared, with public partners typically assuming a disproportionate burden while private partners capture the operational benefits (Macdonald & Cheong, 2014).

Two qualifications are necessary before drawing strong conclusions from this governance literature. First, Pratt (2017) shows that what appears as institutional resilience at the aggregate level can conceal a redistribution of vulnerability toward precarious workers and unpaid volunteers within the system. In other words, resilience achieved through diversification carries costs. Second, the form and scale of support packages in any given crisis also reflect the speed and ideological orientation of the political response. These appear only loosely related to the governance architecture examined above, as the divergence of 2020 responses across similar European systems suggests (Betzler et al., 2020).

Taken together, this literature establishes that funding diversification, governance structure and the design of policy responses all condition how financial stress is transmitted into cultural budget outcomes. It remains, however, largely descriptive and case-specific, and cannot establish whether institutional factors systematically moderate the relationship between financial stress intensity and cultural spending across a broad cross-section of countries when a uniform continuous measure of that stress is applied.

2.3.4. Systemic financial stress and cultural expenditure: towards a continuous measure

The determinants literature of Section 2.2.3 already models cultural expenditure as a continuous function of income, unemployment and the fiscal cycle. What remains discrete in that work is the crisis itself, captured through a binary pre/post-2008 split, a dummy for EU-IMF programmes, or a categorical distinction between episodes. Such indicators register that a country was in crisis, but neither how deeply nor for how long, and they impose the researcher's own periodisation on a phenomenon whose intensity in fact varies continuously. To the best of my knowledge, no study of European cultural expenditure has substituted a market-determined composite index of financial stress for that categorical treatment. The claim here concerns the measurement of the shock, not the use of continuous regressors as such

The Composite Indicator of Systemic Stress (CISS) developed by Holló et al. (2012) offers an alternative. Instead of relying on a qualitative classification of crisis periods, the CISS aggregates fifteen individual stress measures drawn from five market segments, namely money markets, bond markets, equity markets, financial intermediaries and foreign exchange markets, into a single continuous index. Its key methodological innovation lies in the use of time-varying cross-correlations between these segments in the aggregation process. This gives relatively more weight to episodes in which stress is simultaneously present across several market segments, capturing the idea that systemic stress, when it spreads across the financial system as a whole, is more dangerous than an equivalent level of stress confined to a single segment. This property makes the CISS particularly well suited to the object of this thesis. In practice, the empirical analysis relies on the Sovereign CISS (SovCISS), a country-level variant of the CISS that measures stress specifically in sovereign debt markets, available for all eleven countries in the sample (Garcia-de-Andoain & Kremer, 2018). This choice is theoretically consistent with the transmission mechanism outlined in Section 2.3.1, in which sovereign stress specifically constrains government budgets and triggers fiscal consolidation leading to cultural budget cuts. The choice carries a cost, the index is rank-based within each country, so it orders episodes within a national history rather than ranking severity across countries (Section 3.1). As a robustness check, the full country-level CISS is used for the ten countries for which it is available, Greece being the sole exception due to data unavailability (Section 3.7).

Dummies can establish that budgets fell after a date the researcher selected. On the other hand, only a continuous index can ask whether they fall in proportion to stress, only above a threshold, or with what temporal profile. Since simultaneous stress compounds, this predicts potential non-linearities or threshold effects that a simple before/after comparison would obscure. Moreover, because cultural budget cycles are slow-moving, using the annual mean of CISS as a continuous independent variable makes it possible to test whether cultural budgets are more sensitive to the duration of stress than to its peak intensity. The 2008 banking crisis and the 2011-2012 sovereign debt crisis together formed a prolonged, multi-year period of elevated stress, whereas COVID-19 triggered drastic measures within weeks and was met, in several European countries, by rapid and substantial policy support (Betzler et al., 2020). To fully capture this temporal dimension, the average and maximum annual values of the SovCISS are used in the empirical section (Section 3.1).

A continuous measure also allows institutional capacity to be tested against a measure of stress whose intensity and timing are independently determined by financial markets themselves.

3. Empirical analysis

3.1. Data and baseline specification

The empirical analysis draws on a balanced panel of 11 eurozone countries: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Portugal and Spain, observed annually over the period 2006-2024. This sample restricts the analysis to eurozone member states, excluding countries such as the Czech Republic, Denmark, Hungary, Poland, Sweden and the United Kingdom despite the availability of SovCISS data for several of them. The reason is that the transmission mechanism tested here, from banking stress to sovereign stress to fiscal consolidation, presupposes a shared monetary framework and the absence of independent exchange-rate adjustment as a shock absorber, a channel available to non-eurozone states that would confound the interpretation of the coefficient of interest. This restriction is tested rather than assumed, Section 3.7 reports the corresponding results on a panel extended to the EU member states outside the monetary union for which the ECB publishes the SovCISS. Despite being in the eurozone, nor is Cyprus included in this analysis because the ECB does not publish a SovCISS for it.

The dependent variables are cultural expenditure as a share of GDP and cultural expenditure as a share of total government expenditure, both drawn from Eurostat's COFOG (Classification of the Functions of Government) database under the heading "Cultural Services" (GF0802), thereby excluding "recreation" (GF0801), "broadcasting" (GF0803) and "religious and community services" (GF0804). Cultural expenditure in current euros is also collected. It is used descriptively and, deflated by the HICP, as the dependent variable of the per-capita specification of Section 3.7. The panel specification itself requires expenditure to be expressed relative to a scale variable to be comparable across countries of very different economic size.

The central explanatory variable is SovCISS, the sovereign-specific variant of the ECB's Composite Indicator of Systemic Stress (Section 2.3.4), obtained from the ECB Statistical Data Warehouse and aggregated from its native higher-frequency form into annual mean and annual maximum values. The former capturing average yearly stress exposure and the latter the intensity of the most acute episode within the year. The annual mean is used as the primary measure and the maximum as the robustness variant, as slow-moving budget cycles should respond to the duration of stress rather than to its peak.

Institutional quality, used to test H3, is measured by the World Bank's Worldwide Governance Indicators, specifically the Government Effectiveness estimate (WGI-GE; Kaufmann et al., 2011), on a scale from -2.5 to 2.5. Three macroeconomic control variables are included throughout: real

GDP growth, general government debt as a share of GDP, and the unemployment rate, all sourced from Eurostat. Inflation and total fiscal revenue were considered at an earlier stage but dropped from the final specification, on conceptual grounds in both cases. Inflation's price information is already channelled through the HICP-based deflator used to construct the real per-capita measure of Section 3.7, making a separate control here redundant with that later use. Total fiscal revenue, itself a GDP-normalised fiscal ratio, is exposed to the same cyclical denominator effects documented for the dependent variables in Section 3.7, and would enter the specification as a mechanically-driven predictor rather than as a clean confound.

The baseline specification is estimated as:

$$\text{CulturalExp}_{it} = \alpha + \beta_1\,\text{SovCISS}_{it} + \beta_2\,\text{GDPGrowth}_{it} + \beta_3\,\text{Debt}_{it} + \beta_4\,\text{Unemployment}_{it} + \mu_i + \lambda_t + \varepsilon_{it}$$

where CulturalExp_{it} denotes, alternately, cultural expenditure as a share of GDP (testing H1) and as a share of total expenditure (testing H2) for country i in year t; μ_i and λ_t are country and year fixed effects respectively; and ε_{it} is an idiosyncratic error term. This baseline equation is subsequently extended for each remaining hypothesis, an interaction term $\text{SovCISS}_{it} \times \text{WGI-GE}_{it}$ for H3, a lagged stress term $\text{SovCISS}_{i,t-k}$ for H4, and a quadratic term SovCISS^2_{it} for H5, as detailed in the corresponding sections below.

Country fixed effects absorb all cross-country heterogeneity that is constant over the sample window: institutional traditions, the level of the cultural sector, and structural features such as the centralisation of cultural governance discussed in Section 2.2.1. No such level difference, however large, can bias the estimate of β_1, since the within transformation removes it by construction. What fixed effects cannot rule out is a change in these features within a country over the eighteen intervening years that happens to be correlated with SovCISS itself. That possibility is not assumed away here, and is addressed directly through the placebo test and Monte Carlo simulation of Section 3.7. Year fixed effects absorb shocks common to all eurozone countries in a given year, such as the COVID-19 pandemic or eurozone-wide monetary policy decisions, so that β_1 is identified from each country's stress relative to the eurozone average in that year rather than from the level or common trend of stress itself. Standard errors are clustered at the country level throughout, allowing for arbitrary serial correlation within each country's error term, a necessary correction given that SovCISS is itself a persistent, slow-moving series whose shocks are unlikely to be independent across consecutive years

Before turning to what variation identifies β_1, one property of the stress index itself bears on how that variation should be read. The SovCISS is constructed by applying the probability integral transform to each of its six raw components, every observation being replaced by its value under the empirical cumulative distribution function of that series, computed country by country and over an expanding real-time window beginning in December 2006 (Garcia-de-Andoain & Kremer, 2018). The transformation is what makes six heterogeneous inputs commensurable enough to be aggregated, but its authors are explicit that it comes at the cost of the cardinal information contained in the original scale. The implication for a cross-country panel is that a given value of SovCISS locates a country-year within that country's own distribution of sovereign stress, not a scale shared across all eleven. To give an example, a reading of 0.6 in Austria and a reading of 0.6 in Greece denote the same position in two very different national histories, not the same level of market pressure. Within-country identification is unaffected, since country fixed effects absorb the level of each national distribution and the index remains monotone in stress within a country. What the property qualifies is the cross-country comparison on which the residual variation described below relies. It is returned to in Section 3.6, where it accounts for the composition of the extreme tail, in Section 3.7, where it bears on the comparison with discrete crisis indicators, and in Chapter 4.

What variation then identifies β_1 is worth making explicit. Country and year fixed effects together remove 58.9% of the variance of SovCISS prior to estimation (residual SD 0.135, against 0.211 in raw terms), but the two contributions are markedly unequal, year effects account for 48.7 of those percentage points and country effects for only 10.2. Most of the movement in sovereign stress is thus common to the eurozone in a given year, the 2008 and 2011-2012 spikes raised the index across all eleven countries at once while persistent level differences between countries explain comparatively little. Identification therefore rests on the residual within-year differential exposure of countries once that common cycle is absorbed, in practice predominantly the gap between programme and non-programme countries during 2010-2013. This anticipates the low within-R^2 reported below. With nearly half the variance of the stress index absorbed as a common annual cycle, little residual variation remains from which a within-country effect could be identified. Given the rank-based construction just described, that differential should be read as a difference in how far each country stood from its own normal conditions, not as a difference in the absolute severity of the pressure it faced.

Table 1 reports descriptive statistics for the dependant variables, the stress measure and the macroeconomic controls. Cultural expenditure averages 0.42% of GDP and 0.87% of total expenditure across the panel, while SovCISS ranges from 0.029 to 0.946 (median 0.125), reflecting the concentration of acute stress in a small number of country-years.

Prior to estimation, the four time-varying regressors are checked for multicollinearity. Pairwise correlations remain moderate, the strongest being between public debt and unemployment ($r = 0.59$; Table 2a). Variance inflation factors (VIF) confirm the absence of a problematic degree of collinearity, with all values well below conventional thresholds usually fixed at 5 (max VIF = 1.79, for unemployment; Table 2b).

GDP growth, public debt and unemployment are plausible mediators, not pure confounders, of the transmission mechanism, from financial stress to cultural cuts, described in Section 2.3.1. To make this distinction explicit, results for H1 and H2 are reported sequentially in Table 3a: M1 (SovCISS alone, capturing the total effect), M2 (adding GDP growth), and M3 (adding public debt and unemployment, the baseline specification retained throughout the remainder of this chapter). The contrast between β_1 in M1 and M3 shows how the estimated association changes as these macro-fiscal mediators are introduced.

Two panel diagnostics inform the choice of estimator: a Pesaran (2004) test for cross-sectional dependence (CD) and a Wooldridge (2010) test for serial correlation. Both reject their nulls, as reported in Section 3.3, and Driscoll-Kraay standard errors (Driscoll & Kraay, 1998) are therefore reported alongside country-clustered errors throughout.

The small number of clusters requires a further check. The wild cluster bootstrap of Cameron et al. (2008) is used on every reported null. With G = 11, the Rademacher weight space contains only 2^{11} = 2048 distinct sign vectors, so the bootstrap, imposing the null on the restricted residuals, is computed by full enumeration instead of simulation. This yields exact and deterministic p-values (MacKinnon et al., 2023), computed on the tail of the bootstrap distribution corresponding to each hypothesis's stated alternative. The confidence intervals reported throughout are the analytic clustered ones, read descriptively rather than as tests (Section 3.2). All models are estimated in R with the fixest package (Bergé, 2018), for its handling of high-dimensional fixed effects and multi-way clustering.

3.2. Hypotheses

The review conducted converges on five testable hypotheses, which the empirical analysis addresses in turn.

Four of the five hypotheses are directional and are tested one-sided: H1, H2 and H4 against a negative alternative, H3 against a positive one. H5 is the only two-sided hypothesis, since the threshold logic of Section 2.3.4 predicts a departure from linearity without specifying its direction. All p-values reported in the text follow the alternative stated for each hypothesis, their two-sided counterparts are given in the Annex. Confidence intervals, by contrast, are reported two-sided throughout, as a descriptive measure of the range of effects compatible with the data rather than as a test.

The first hypothesis concerns the direct effect of financial stress on public cultural budgets. Because cultural spending behaves as a superior good and falls within the most exposed category of discretionary expenditure, it is expected to contract more than proportionally as financial stress intensifies. Higher systemic financial stress, as measured by SovCISS, should therefore reduce public cultural expenditure as a share of GDP (H1).

$$H_0 : \beta_{\text{SovCISS}} = 0 \qquad ; \qquad H_A : \beta_{\text{SovCISS}} < 0$$

H2 isolates a compositional effect distinct from the first. Even net of the aggregate contraction of public budgets, culture should lose ground to other items, consistent with the political economy of austerity set out in Section 2.3.1, where entitlement programmes are shielded and discretionary items bear a disproportionate share of the adjustment. Stress reduces the share of culture within total public expenditure, independently of its effect on the overall envelope.

$$H_0 : \beta'_{\text{SovCISS}} = 0 \qquad ; \qquad H_A : \beta'_{\text{SovCISS}} < 0$$

H3 predicts a weaker effect where institutional quality is higher. The resilience argument of Section 2.3.3 holds that stronger institutional capacity limits the transmission of a fiscal shock into cultural budget cuts. That buffering may itself be overwhelmed when stress is generalised enough to reach all tiers of government at once. The moderator is the World Bank's Government Effectiveness indicator (WGI-GE) in preference to decentralisation directly (Chapter 4).

$$H_0 : \gamma_3 = 0 \qquad ; \qquad H_A : \gamma_3 > 0, \quad \text{where } \gamma_3 \text{ denotes the coefficient on the interaction term } (\text{SovCISS}_{it} \times \text{WGI-GE}_{it}).$$

The fourth hypothesis addresses the persistence of cuts documented empirically across Europe (Almeda et al., 2018) and interpreted through the reversed ratchet effect (Jaén-García, 2021), developed in Section 2.2.3. The effect of financial stress on cultural expenditure is therefore expected to persist beyond the period of peak stress itself (H4).

$$H_0 : \gamma_4 = 0 \qquad ; \qquad H_A : \gamma_4 < 0, \quad \text{where } \gamma_4 \text{ denotes the coefficient on } \text{SovCISS}_{i,t-k}.$$

H5 allows the relationship to be non-linear. This follows from the threshold logic built into the CISS methodology itself (Section 2.3.4), where systemic risk compounds once stress spreads across market segments simultaneously. It is also consistent with the observation that severely stressed states, notably those under EU-IMF programmes, took disproportionately deeper cuts than moderately stressed ones (Srakar et al., 2017).

$$H_0 : \gamma_5 = 0 \qquad ; \qquad H_A : \gamma_5 \neq 0, \quad \text{where } \gamma_5 \text{ denotes the coefficient on the quadratic term } \text{SovCISS}^2_{it}.$$

Three hypotheses formulated during the design were not carried into estimation. The first concerned private substitution for public withdrawal, following the crowding-in/crowding-out debate of Section 2.2.2. No continuous cross-country panel documents private cultural expenditure in Europe, so it could not be tested at all. Another distinguished a Northern from a Southern European response, motivated by the asymmetry Rubio-Arostegui and Rius-Ulldemolins (2020) document. A geographical dichotomy is only an approximate indicator of the institutional capacity it is meant to represent, whereas H3 measures that capacity directly. The geographic partition is not the programme partition examined in Section 3.4. The third hypothesis concerned the composition of adjustment within cultural budgets, where the political economy of consolidation predicts that gross fixed capital formation absorbs more of the cut than current spending (Breunig & Busemeyer, 2012). This last one was abandoned on a power criterion, the minimum detectable effect (MDE; Bloom, 1995) on the coefficient of interest exceeded plausible economic magnitudes by two orders of magnitude. That calculation uses the standard error and the sample alone. It was made without reference to the sign or the significance of the coefficient.

3.3. The direct effect of financial stress on cultural expenditure

The baseline specification described in Section 3.1 is estimated separately for cultural expenditure as a share of GDP (H1) and as a share of total government expenditure (H2), using SovCISS annual mean as the primary stress measure. Prior to interpreting the coefficients, two diagnostic tests inform the choice of inference. The Pesaran test rejects the null of cross-sectional independence (z = -2.72, p = 0.007), but the average pairwise residual correlation is -0.084, close to the mechanical benchmark of $\frac{-1}{N-1}$ = -0.100 implied by within-year demeaning. The rejection therefore reflects the test's sensitivity to a small residual beyond the mechanical component, but this rejection is clear and Driscoll-Kraay standard errors are reported accordingly.

The Wooldridge test rejects the null of no first-order serial correlation (F = 377.42, p < 0.001); the estimated autocorrelation of the two-way residuals is $\rho = 0.795$, confirming that country-clustered standard errors, which allow for arbitrary within-country serial correlation, are the appropriate baseline. Driscoll-Kraay standard errors, robust to both cross-sectional dependence and serial correlation, are reported alongside country-clustered errors in Table 3a. We can further note that sensitivity of β_1's Driscoll-Kraay standard error to the truncation lag is negligible: 0.026, 0.024 and 0.021 for lags one, two and three respectively. The reported specification uses DK(2).

Two estimands are in play, and the distinction governs how Table 3a is read. M1 identifies the association before the macro-fiscal channels enter whereas M3 identifies what remains once growth, debt and unemployment are conditioned on. Neither is a clean total effect. M3 conditions on channels through which the hypothesised mechanism operates, while M1 leaves the association open to confounding by public debt, which plausibly causes sovereign stress through the spread channel as much as it transmits it. The sequence is reported in full for that reason.

Table 3a reports the baseline estimates. The coefficient on SovCISS is negative in both specifications, consistent in sign with H1 and H2, but far from statistically significant. In the baseline specification (M3), β_1 = -0.025 (SE = 0.062) for cultural expenditure as a share of GDP, and β_1 = -0.036 (SE = 0.110) for its share of total expenditure. Adding the macro-fiscal controls moves the two coefficients in opposite directions. For the GDP share, β_1 is essentially unchanged across the three columns (M1 = -0.013, M2 = -0.024, M3 = -0.025). The raw association is already close to zero, and the controls neither absorb nor create one. For the share of total expenditure, by contrast, the raw association is both larger and statistically significant, SovCISS alone yields M1 = -0.123 (SE = 0.048), significant at the 5% level, and this coefficient is progressively absorbed as GDP growth, public debt and unemployment enter (M2 = -0.105, M3 = -0.036), until it is no longer distinguishable from zero. The sequence is read here as mediation, not as spurious correlation. The distinction is not statistical, since the same three columns are consistent with either reading. It rests on the causal ordering of Section 2.3.1, where growth, debt and unemployment are the channels through which sovereign stress reaches discretionary budgets. Their absorption of the coefficient is what the mechanism predicts. H2, as formulated in Section 3.2, concerns the compositional margin net of the aggregate envelope, and M3 is the specification frozen against that formulation. It is reported as the specified test and is not redesignated after the fact. The normalisation of the dependent variable by GDP is discussed in Section 3.7.

Neither estimate approaches conventional significance thresholds, a result reported in Table 3b (p = 0.385 and p = 0.364 respectively), for the M3 specifications computed by full enumeration of the 2 048 Rademacher vectors. None of the macroeconomic controls, GDP growth, public debt, or unemployment, reach significance either.

This absence of a detectable direct effect should be read against the very low within-R^2 of the two models (0.014 and 0.117 respectively). It indicates that, once country and year fixed effects are accounted for, the remaining within-country variation explained by SovCISS and the controls is limited. Most of the explanatory power of the model (overall R^2 above 0.90) is absorbed by the fixed effects themselves, not by the time-varying regressors of interest. This pattern qualifies the cross-country evidence of Section 2.3.2. Srakar et al. (2017) found a pooled crisis effect visible only once cross-country heterogeneity is modelled explicitly. The present specification absorbs that heterogeneity through country fixed effects, and may therefore identify β_1 from comparatively little residual variation.

Whether this null result reflects a genuine absence of a linear, contemporaneous relationship, as opposed to a lagged (H4) or non-linear (H5) one, is examined in the sections 3.5 and 3.6 respectively. The interpretive weight of these null results depends on what the design could have detected. Each null is therefore reported alongside its MDE, the smallest true coefficient the design would have had an 80% probability of detecting as significant at the 5% level under the directional alternative, given the realised precision of the estimate. For H1, with a clustered SE of 0.062 and critical values based on G-1 = 10 degrees of freedom, the MDE is 0.17 percentage points of GDP per unit of SovCISS. The equivalent MDE for H2 is 0.30 percentage points of the total-expenditure share per unit of SovCISS, likewise far beyond its point estimate of -0.036, so that the same interpretation extends to the compositional specification.

Translated into a representative stress episode (ΔSovCISS ≈ 0.55, the median within-country amplitude in the sample), the design can only detect stress-induced changes exceeding roughly 21% of the average cultural budget. The 95% confidence interval on β_1 is compatible with anything from a contraction of 20.9% to an expansion of 14.5% of that same budget in response to a major stress episode. The bound can be placed against the largest observed contraction in the panel. Greek cultural expenditure fell 54% between 2008 and 2010 (Section 3.8), a contraction the design would have detected, though under a stress movement well above the median episode used here. A cut a third as deep would have passed unnoticed at any dose this panel contains. The nulls of this section therefore exclude responses of Greek magnitude across the panel, they are not evidence that financial stress leaves cultural budgets unaffected.

3.4. Institutional quality as a moderator of financial stress

Section 3.2 formulated H3 around the resilience argument developed in Section 2.3.3. This hypothesis is tested by extending the baseline equation with an interaction term between SovCISS and institutional quality, together with the latter's own main effect:

$$\text{CulturalExp}_{it} = \alpha + \beta_1\,\text{SovCISS}_{it} + \beta_2\,\text{GDPGrowth}_{it} + \beta_3\,\text{Debt}_{it} + \beta_4\,\text{Unemployment}_{it} + \delta\,\text{WGI-GE}_{it} + \gamma_3\,(\text{SovCISS}_{it} \times \text{WGI-GE}_{it}) + \mu_i + \lambda_t + \varepsilon_{it}\,.$$

Both SovCISS and WGI-GE are mean-centred prior to interaction, so that β_1 retains its interpretation as the effect of financial stress at average institutional quality. Unlike H1 and H2, where the regressor of interest is SovCISS itself, the relevant robustness concern here is the stability of γ_3, the interaction coefficient, as the moderator's own main effect and macroeconomic controls are added sequentially (Table 4a). γ_3 remains between 0.042 and 0.048 across the sequence (Table 4a), indicating that the interaction is not an artefact of a particular control configuration.

Table 4a reports the baseline estimate alongside Driscoll-Kraay standard errors. The interaction coefficient is small and statistically indistinguishable from zero (γ_3 = 0.045, country-clustered SE = 0.041; Driscoll-Kraay SE = 0.056), so H_0 cannot be rejected. Its sign is nonetheless consistent with H_A, the point estimate implies that higher institutional quality doesn't amplify but absorbs the negative effect of sovereign stress. The marginal effect of stress rises monotonically across the distribution of institutional quality, from -0.057 at the 10th percentile to 0.011 at the 90th (Table 4b). None of these conditional effects is individually significant, but their monotonic ordering across the distribution of institutional quality traces exactly the pattern the resilience literature predicts (Getzner, 2015; Rizzo & Throsby, 2006), even though the sample cannot establish it with statistical confidence. A wild cluster bootstrap on the interaction coefficient (Cameron et al., 2008) confirms the country-clustered inference (p = 0.231, computed by full enumeration; see Table 3b).

The MDE on γ_3 at 80% power is 0.109, roughly two and a half times the point estimate, so that the moderation null bounds the interaction away from large effects without establishing its absence.

A further concern is that the WGI-GE interaction reproduces the effect of EU-IMF programme status, which is where the identifying variation of Section 3.6 concentrates. Table 4c addresses this by adding an interaction between SovCISS and an indicator for the three programme countries, i.e. Greece, Ireland and Portugal (Section 2.3.2). Note that this is not the geographic partition set aside in Section 3.2. Ireland is a programme country outside the Southern group, and Spain and Italy are Southern countries that never entered a programme. Adding the programme interaction raises γ_3 by half, from 0.045 to 0.069 (SE = 0.039), while the programme term is itself of comparable magnitude (0.066, SE = 0.050), the within-R^2 rises from 0.072 to 0.089. Under the enumerated bootstrap the interaction reaches p = 0.052 against its stated one-sided alternative, while the programme term does not (p = 0.329, two-sided, no direction being predicted for it). The interaction is therefore not displaced by programme status; it does not follow that the two capture unrelated variation. Their correlation over the sample is -0.447, and when two correlated regressors enter jointly, the estimator cannot cleanly attribute their shared variance to one rather than the other, standard errors inflate and the split between them is partly arbitrary. Table 4d reports the marginal effect implied for each country at its own mean of institutional quality, none is individually significant.

The main effect of WGI-GE is negative and not distinguishable from zero under country-clustered inference (δ = -0.070, SE = 0.043), though it reaches significance under Driscoll-Kraay errors; it is not interpreted here. No ex-ante hypothesis concerned the level effect of institutional quality, and the within-country variation identifying it is small relative to the indicator's published margins of error (Kaufmann et al., 2011).

Taken together, these results offer no confirmation of H3 on the specification against which it was frozen; the variant that comes closest is not redesignated as the test. They are directionally consistent with the theoretical mechanism motivating it, robust to the choice of macroeconomic controls, and not reducible to programme status. Chapter 4 returns to the measurement compromise involved in proxying decentralisation through WGI-GE. A COFOG-based decentralisation index was constructed as an alternative but could not be retained, as Eurostat does not report the required sub-national breakdown for Germany and Austria, the sample's two paradigmatic federal states.

3.5. Persistence of financial stress effects

The fourth hypothesis addresses the temporal dimension of the stress-expenditure relationship. The empirical literature documents that cultural budget cuts made under fiscal duress tend to become embedded in subsequent budgetary baselines, not reversed once stress subsides (Almeda et al., 2018; Jaén-García, 2021), suggesting that the effect of financial stress on cultural expenditure should persist beyond the period of peak stress itself. To test this, the baseline specification of Section 3.1 is extended with the first lag of the stress variable:

$$\text{CulturalExp}_{it} = \alpha + \beta_1\,\text{SovCISS}_{it} + \gamma_4\,\text{SovCISS}_{i,t-1} + \beta_2\,\text{GDPGrowth}_{it} + \beta_3\,\text{Debt}_{it} + \beta_4\,\text{Unemployment}_{it} + \mu_i + \lambda_t + \varepsilon_{it}\,.$$

where γ_4 captures the effect of past sovereign stress conditional on contemporaneous stress. The formal test is $\text{H}_0 : \gamma_4 = 0$ against $\text{H}_\text{A} : \gamma_4 < 0$. The lag drops the 2006 cross-section, leaving 198 observations over 2007-2024. Re-estimated on that sample, the contemporaneous-only benchmark yields -0.029 (clustered SE 0.060), confirming that this loss does not alter the picture (Table 5).

Because H4 concerns whether the effect of stress outlasts the stress itself, the quantity of interest is the cumulative two-year response, $\beta_1 + \gamma_4$, tested jointly. The two terms are correlated in the pooled data (r = 0.702), a consequence of stress episodes spanning several years, but the year fixed effects absorb most of that common component, the standard error on is 0.0598 when the β_1 contemporaneous term enters alone and 0.0587 when the lag is added. Precision is therefore not the constraint. The individual coefficients are reported in Table 5 and suggest the cumulative effect is simply the quantity H4 concerns.

The results provide no support for H4. The coefficient on lagged stress is γ_4 = -0.026, negative as the hypothesis predicts (Table 5), and the wild cluster bootstrap yields p = 0.346 by full enumeration (Table 3b). The cumulative two-year effect, computed by the delta method with country-clustered variance, is -0.043 (95% CI [-0.182; 0.095]; Table 5), and a Wald test of the joint nullity of both stress terms yields F(2,10) = 0.23, p = 0.800. Its confidence interval implies that contemporaneous standard error cannot be distinguished from the cumulative two-year coefficient's one, so aggregating over two years buys no precision. Both terms carry the predicted sign when they enter jointly (β_1 = -0.017, γ_4 = -0.026), but neither is estimated precisely enough for that to count: adding the lag lifts the within-R^2 from 0.015 to 0.018, and the MDE on γ_4 at 80% power is 0.123, roughly five times the point estimate. The cumulative response runs into the same limit as the contemporaneous one.

Conditional on contemporaneous stress, there is no detectable transmission of sovereign stress to cultural expenditure at a one-year horizon. This does not rule out the entrenchment documented in the literature, which is a property of the persistence of the expenditure series rather than of a delayed stress effect and would require a dynamic specification to estimate directly. Chapter 4 sets out why that specification was unavailable here. Two features of the results nonetheless bear on it. The very low within-R^2 across all specifications indicates that the temporal variation of cultural expenditure within countries is largely disconnected from annual macro-financial fluctuations, consistent with a viscous budget series in which level-persistence remains plausible absent a lagged stress effect. What's more, the stability of the null across clustered, Driscoll-Kraay and bootstrap inference shows it is not an artefact of the treatment of the error structure.

Extending the specification to a second lag confirms this. On the reduced sample of 187 observations (2008-2024), the first lag alone reaches p = 0.034 against the stated alternative, but the two-lag specification was not pre-specified, extending it further in response to this marginal result would constitute specification search. Plus, its coefficient (-0.041) lies within the confidence interval of the pre-specified model and below the MDE of 0.123. It does not constitute evidence of persistence the pre-specified design could not detect. The quantities H4 concerns, i.e. the cumulative effect (-0.074, 95% CI [-0.241; 0.093]) and the joint test (F(3,10) = 1.58, p = 0.255), remain far from any threshold. Past sovereign stress adds no detectable explanatory power for current cultural expenditure at either horizon.

With Sections 3.3 and 3.4 this completes a consistent pattern, no uniform linear effect of sovereign stress on cultural expenditure, contemporaneous or delayed. If the relationship documented in the discrete-crisis literature exists, it is likely concentrated in episodes of extreme stress (Srakar et al., 2017).

3.6. Non-linearity of the stress-expenditure relationship

The fifth hypothesis addresses the functional form of the relationship (H5). Following the threshold logic embedded in the CISS methodology itself (Section 2.3.4), the baseline specification of Section 3.1 is extended with a quadratic stress term:

$$\text{CulturalExp}_{it} = \alpha + \beta_1\,\text{SovCISS}_{it} + \gamma_5\,\text{SovCISS}_{it}^2 + \beta_2\,\text{GDPGrowth}_{it} + \beta_3\,\text{Debt}_{it} + \beta_4\,\text{Unemployment}_{it} + \mu_i + \lambda_t + \varepsilon_{it}\,.$$

where SovCISS is mean-centred prior to squaring, so that β_1 retains its interpretation as the marginal effect of stress at the sample mean, not at zero. The formal test is $\text{H}_0 : \gamma_5 = 0$ against $\text{H}_\text{A} : \gamma_5 \neq 0$. H5 is the only hypothesis of the set formulated as two-sided, since the threshold logic predicts a departure from linearity without imposing its direction ex ante. The p-values and the MDE reported below are accordingly two-sided. One diagnostic conditions the reading of the estimates. Because the distribution of SovCISS is strongly right-skewed (range = [0.029; 0.946], median = 0.125), centring only partially decorrelates the linear and quadratic terms, whose pooled correlation remains 0.775. The two coefficients thus share part of a weak common signal, their individual magnitudes are unstable across specifications. What can be read is the curvature test on γ_5 and the marginal effects of stress evaluated across the distribution, rather than either coefficient read in isolation.

Table 6a reports the estimates, with the linear model of Section 3.3 re-displayed as the nested benchmark. The results provide no support for H5. The coefficient on the quadratic term is γ_5 =

-0.080 (clustered SE = 0.117; Driscoll-Kraay SE = 0.095), statistically indistinguishable from zero under every inference method employed. The cluster-robust Wald test yields $p = 0.512$, and the wild cluster bootstrap yields $p = 0.618$ (Table 3b). The within-R^2 rises only from 0.014 to 0.020 with the inclusion of the quadratic term. Following the MDE framework, the design could only have detected a curvature of approximately $|\gamma_5| \approx 0.363$ at 80% power, more than four times the estimated magnitude, so that this null should again be read as the absence of a detectable non-linearity rather than as affirmative evidence of linearity. The null hypothesis is not rejected.

As with H3 and H4, the point estimates are nonetheless directionally consistent with the theoretical mechanism motivating the hypothesis. The estimated curvature is concave, implying that the marginal effect of stress declines monotonically across the distribution of SovCISS, from 0.037 at the 10th percentile to -0.036 at the 90th, crossing zero between the median and the 75th (Table 6b). None is individually significant. The ordering is the one the compounding logic of systemic stress predicts, as the ordering of Section 3.4 was the one the resilience literature predicts, in neither case established statistically.

A quadratic polynomial is, moreover, only one representation of the threshold logic. Programme entry suggests a regime change beyond a critical stress level, not a smooth curvature. As a robustness check, the quadratic term is therefore replaced by a linear spline with a single knot, placed alternately at the 75th and at the 90th percentile of SovCISS, so that the coefficient on the spline term measures the additional slope prevailing beyond the threshold (Table 6a). The spline was not specified ex-ante, so the quadratic stands as the test of H5 and the spline as a check on it. The slope supplement is negative at both knots but imprecisely estimated. Beyond the 90th percentile the total marginal effect of stress is -0.080, with a delta-method confidence interval from -0.275 to 0.115 (Table 6a). Neither representation of the threshold logic is detectable at this level of precision. The MDE on γ_5 is 0.363, and the spline supplements are estimated with confidence intervals spanning both signs. The result is unchanged when the dependent variable is replaced by cultural expenditure as a share of total government expenditure ($\gamma_5 = -0.170$, SE = 0.255), and the Driscoll-Kraay standard errors are stable across truncation lags (0.092, 0.095 and 0.097 for lags one to three).

One final diagnostic delimits what this specification can and cannot identify. Most of the twenty-one observations lying above the 90th percentile of SovCISS are concentrated in three countries, fourteen observations in the programme group, seven of them for Greece only. Each country's index reaches high values during its own worst episode whatever the absolute severity of that episode, so the upper decile of the pooled distribution mixes country-years of genuinely extreme pressure with country-years that were merely unusual by national standards. Identifying a non-linearity in SovCISS therefore rests on a tail whose most influential observations are the same country-years that would identify an EU-IMF programme effect. A leave-one-country-out (LOO) sensitivity check confirms the implication. Excluding any single country other than Greece leaves γ_5 within a range of -0.184 (Portugal) to -0.010 (Spain) around the full-sample estimate, but excluding Greece reverses its sign (γ_5 = 0.031), indicating that the weak concave pattern documented above is carried by the Greek observations alone. This should be read as a limit of identification specific to a sample in which the tail of the stress distribution is populated by a small number of country-year observations, rather than as a specification failure. Within this panel, a genuine curvature in the stress-expenditure relationship and a response specific to programme countries are observationally equivalent, and the data cannot separate them.

Taken together, the results of Sections 3.3 to 3.6 establish a consistent pattern. Sovereign financial stress exhibits no detectable effect on public cultural expenditure in this eurozone panel, whether the relationship is specified as linear, lagged, moderated by institutional quality, or non-linear. Section 3.7 subjects these baseline nulls to a systematic robustness battery before Section 3.8 draws together what this five-hypothesis pattern implies for the question posed at the outset.

3.7. Robustness

The null results established in Sections 3.3 to 3.6 carry interpretive weight only to the extent that they are not artefacts of the choices embedded in the baseline design: the measure of financial stress, the normalisation of the dependent variable, the composition of the sample, and the treatment of the error structure. A battery assembled after inspecting which variants behave favourably would be specification search. This is not the case of the placebo, estimator-comparison and extended-panel exercises that follow, which were added after estimation in response to specific objections. Two corrections already applied throughout, the enumerated wild cluster bootstrap and the Driscoll-Kraay standard errors motivated by the Pesaran (2004) CD test rejection of Section 3.3, are not repeated here. The full battery appears in Table 7a for the GDP share and Table 7b for its share of total expenditure, the discrete-crisis comparison appears in Table 8.

The identification strategy of Section 3.1 assumes that contemporaneous stress is not itself a response to anticipated future adjustments in cultural expenditure. A placebo test extended the baseline with the one-and-two-year-ahead value of SovCISS. These leads should carry no explanatory power once contemporaneous stress is controlled for, but the results are different. The one-year lead is small and indistinguishable from zero but adding a second lead strongly rejects the joint nullity of the two leads ($F(2,10)=13.85$; enumerated bootstrap $p = 0.001$). Coefficients in the joint test have opposed signs (-0.055 and 0.081), suggesting that the test captures the dynamic of the stress process more than a budgetary anticipation. The strict exogeneity can therefore not be maintained (Table 7c). Against this, the first-difference and instrumental-variable (Anderson & Hsiao, 1981) checks return -0.029 and -0.032 against a FE estimate of -0.025. These estimators do not share a probability limit once strict exogeneity fails. If endogeneity were substantial, the IV correction should diverge from the FE estimate by the order of the bias. The observed gap is 0.007 against a worst-case bound of 0.142 from the Monte Carlo feedback scenario (Chapter 4) pointing to a realised bias well below the detectable threshold.

A first group of checks replaces the stress measure itself. SovCISS was selected on theoretical grounds, as the variant of the CISS family that most directly captures the sovereign leg of the transmission mechanism (Section 2.3.1). However, this choice involves two discretionary elements: the restriction to sovereign debt markets and the aggregation of the underlying higher-frequency series into an annual mean. The first is relaxed by substituting the full country-level CISS, which aggregates stress across five market segments (Holló et al., 2012). Because the country-level CISS is unavailable for Greece, this substitution mechanically restricts the panel to ten countries, so that a naïve comparison with the eleven-country baseline would confound the change of measure with the change of sample. To separate the two, the SovCISS baseline is first re-estimated on that same ten-country subsample, holding the measure fixed while removing Greece. Dropping Greece alone flips the sign for both dependent variables (β_1 = 0.050, SE = 0.066, for the GDP share; β_1 = 0.103, SE = 0.097, for the expenditure share). These are, by construction, the Greece-exclusion estimates of the jackknife reported below. Only against this benchmark is the measure itself varied. Replacing SovCISS with the country-level CISS on that identical ten-country sample returns the point estimate to the baseline's negative sign for both dependent variables (column R1 of Tables 7a and 7b). Neither perturbation yields an estimate distinguishable from zero, from the bridging estimate, or from the baseline. On a fixed ten-country sample the measure change alone moves the coefficient from 0.050 to -0.031, so the exercise bounds neither a sample effect nor a measure effect. What it establishes is that the baseline sign survives neither of them. This extends to the linear baseline the identification limit already documented for the quadratic term in Section 3.6.

The euro-area aggregate CISS is identical across countries within a given year, hence perfectly collinear with the year fixed effects. It is estimated instead with country fixed effects but omitting year effects (R2). This captures the response to the common component of European financial stress rather than to differential exposure within a given year. This is the only specification in the battery to return a coefficient distinguishable from zero (Table 7a and 7b). The same specification estimated on SovCISS without year effects is itself positive (0.039 and 0.031 for the two dependent variables), the appropriate comparison as both rest on the same identifying variation. The positive sign is thus not specific to the euro-area measure but a structural feature of dropping year fixed effects. With the common component of stress unabsorbed, the coefficient partly reflects the joint upward drift of cultural expenditure and financial stress over the period, not a within-year cross-country response, and cannot be read causally (Section 3.8). Two notable examples of this are the 2009 recovery packages and the 2020 emergency support, both of which raised expenditure precisely when stress peaked. The single significant coefficient in the battery therefore validates the identification choice defended in Section 3.1 rather than contradicting the baseline finding.

Figure 1 places six of the eight battery variants on a common scale. R5, expressed as a semi-elasticity, and R6, a range of eleven estimates, are discussed below. Four further variants perturb the measure, the specification and the error structure: the annual maximum of SovCISS (R3), a ParlGov left-right control (R4, the sample reduced to 207 observations by the uncoded French and Spanish 2024 cabinets), exclusion of the 2020–21 COVID episode (R7, β_1 = -0.026), and two-way country-by-year clustering (R8; Cameron et al., 2011). Together with the checks already discussed, they place seven estimates of β_1 between -0.031 and -0.014, once the two year-fixed-effect-free variants (R2, 0.070 and 0.039) are set aside. That dispersion is one tenth of the minimum detectable effect, and smaller than a tenth of the width of any single confidence interval in the figure. How the model is specified moves β_1 an order of magnitude less than sampling noise within one specification.

Figure 1

Robustness of the baseline estimate of financial stress on cultural expenditure

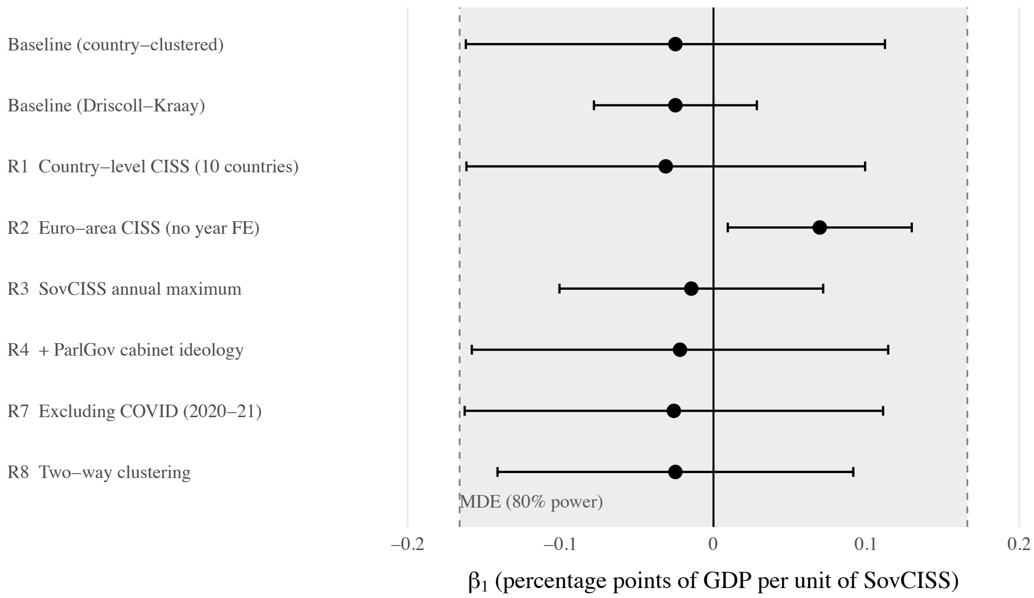

95% CI from country–clustered SE (G–1 = 10 df). Shaded band: MDE = ±0.17 at 80% power (Section 3.3).
R2 variants omit year FE (different identifying variation). R5 (log semi–elasticity) and R6 (jackknife range) omitted: not on the percentage–points scale.

A further check targets the construction of the dependent variable, a concern that is mechanical and not substantive, both dependent variables are ratios whose denominators respond to the cycle. Cultural expenditure as a share of GDP can rise during a crisis simply because nominal GDP falls faster than nominal cultural spending, biasing against detecting H1. Its share of total expenditure can fall simply because automatic stabilisers inflate the denominator, biasing in favour of detecting H2. To remove both channels at once, R5 replaces the dependent variable with the logarithm of real cultural expenditure per capita, deflated by the HICP, and adds the logarithm of real GDP per capita as a control so that the income elasticity documented in Section 2.2.3 is held explicitly rather than absorbed into the ratio; because this is intrinsic to the level of expenditure rather than its budgetary share, it applies to H1 only. Over an episode of median within-country amplitude it implies a contraction of 9.8% of average cultural expenditure against 3.2% in the GDP-share baseline. Both are statistically indistinguishable from zero and from each other. The gap is running in the direction the denominator predicts and the specification could only have detected contractions exceeding 30.8%. R5 identifies a defect in the GDP-share measure on which H1 was pre-registered, not a reason to substitute it after the fact. The ratio is least reliable in exactly the conditions the hypothesis concerns.

The sample is then varied, a leave-one-country-out jackknife (R6) re-estimates the baseline eleven times for each dependent variable, excluding each country in turn. For cultural expenditure as a share of GDP, β_1 ranges from -0.056 (excluding Spain) to 0.050 (excluding Greece); for its share of total expenditure, from -0.112 (excluding Spain) to 0.103 (excluding Greece). In both cases Greece is the single most influential exclusion and the only one that reverses the sign of β_1, extending to the linear baseline the fragility already documented for the quadratic term in Section 3.6.

Two variants carry a reading beyond the null they confirm. The ParlGov control (R4) is itself negative and marginally significant (-0.009 and -0.017 for the two dependent variables, $p < 0.10$, the direction Potrafke (2011) reports for right-leaning cabinets). A time-varying political determinant is detectable where stress is not, without displacing β_1; the baseline null is a bound and not a blind spot. The two-way standard errors (R8), by contrast, fall below their country-clustered counterparts and offer no independent reassurance. With only nineteen year clusters this estimator is known to understate uncertainty; they are read alongside, not in place of, the Driscoll-Kraay correction of Section 3.3.

A complementary check addresses a different question from the eight variations above. It is not whether β_1 is robust to how the design is built, but whether a continuous stress measure adds anything over the discrete crisis indicators common in the literature it seeks to discipline (Section 2.3.4). Three specifications are compared for cultural expenditure as a share of GDP. A first one using the three period dummies employed by Srakar et al. (2017) and related studies (2008 crisis, sovereign debt crisis, COVID), then one using SovCISS alone, and an encompassing specification with both. Because the crisis dummies are purely temporal and therefore collinear with year fixed effects, all three models retain only country fixed effects, replicating the design of the discrete-crisis literature rather than the baseline specification of Section 3.1. The dummies remain jointly significant at 10% level even in the presence of SovCISS ($F(3,10) = 2.85$, $p = 0.091$), while SovCISS adds nothing once the dummies are included ($F(1,10) = 0.29$, $p = 0.603$). The dummy-only specification fits better than SovCISS alone by a wide margin (AIC -684.7 against -673.7) and is not improved by adding SovCISS (-683.6), which is the point (Table 8). The 2008 dummy is itself positive and significant (0.041). It is the same procyclical-denominator artefact documented for R2 and for the Greece exclusion: what a crisis indicator picks up on a GDP-normalised series is partly the collapse of the denominator. The measurement property documented in Section 3.1 supplies a specific reason why this should be so. Because the probability integral transform rescales each country's stress history onto a common unit interval, SovCISS is silent on how much more severe the Greek episode was than the Austrian one, whereas a crisis dummy still separates the country-years the literature treats as crises from those it does not. Where the quantity of interest is the ordering of severity across countries, a rank-preserving transformation discards precisely the information a dummy retains. Taken at face value the result qualifies rather than undermines the case of Section 2.3.4 for a continuous measure. Within this panel, discrete indicators and the tail of SovCISS pick out the same episodes, which does not make the continuous measure redundant in a broader or differently structured sample. This tension was left for Chapter 4.

A last qualification points to a question the battery cannot answer from within the eurozone sample. The exclusion of non-eurozone member states, defended in Section 3.1 on the grounds that exchange-rate adjustment offers these states a shock absorber unavailable inside a monetary union, rests on the transmission mechanism rather than on data availability: the ECB publishes the SovCISS for several EU countries outside the euro area. Extending the panel to Czechia, Denmark, Hungary, Poland and Sweden (sixteen countries, 304 observations) makes the restriction testable rather than assumed (Table 10a).

A fully interacted specification, in which every coefficient may differ across the two groups, cannot reject the common coefficient vector ($F(4, 15) = 0.99$, $p = 0.444$; Table 10b). The interaction on the stress term alone, however, is significant at the 10% level under a Webb-weighted wild cluster bootstrap (0.232, SE 0.149, $p = 0.090$), the Webb weights replacing full enumeration, feasible at 2^{16} vectors but computationally costly. The two results are consistent rather than contradictory. The joint test pools four restrictions, three of which concern controls that plainly do not differ, which dilutes the statistic. The stress response does plausibly differ across monetary regimes. Its direction, however, runs against the mechanism originally invoked. Non-eurozone slopes are on average more negative (-0.046 against 0.019 under the Mean Group estimator (Pesaran & Smith, 1995) and -0.130 against -0.025 in separate estimation), which the exchange-rate absorber argument does not predict. The restriction is vindicated as a modelling choice without being vindicated in its stated rationale, and the reason it is not is that heterogeneity within each group swamps the difference between them.

On the eleven-country panel, replacing the common-slope FE with the Mean Group and CCEMG estimators leaves the coefficient indistinguishable from zero (0.019 and -0.022 against the -0.025 baseline; Table 9), so the null does not stem from imposing a common slope. That slope is nonetheless not common: a poolability test rejects equality of the country-specific slopes ($F = 5.90$, $p < 0.001$; Table 9), and a sequential Kapetanios (2003) procedure rejects global poolability on this sample, excluding Ireland then Portugal (Table 10c, Panel A). Whether that heterogeneity follows the monetary border is a distinct question, and a Chow test on the eurozone/non-eurozone partition answers no ($F(4,15) = 0.99$, $p = 0.444$; Table 10b). Re-run on the sixteen-country panel, the procedure excludes Ireland, Portugal and Hungary, all programme countries. While also being a programme country, Greece stays inside the poolable group (Table 10c, Panel B). The explanation is that the test isolates slope heterogeneity, not the leverage that a few extreme Greek country-years exert in Sections 3.6 and 3.7, and the two need not coincide.

Re-estimating the five hypothesis specifications of this chapter on the extended panel (Table 10d) leaves all coefficients statistically indistinguishable from zero, while four of the five changed their sign. This is the same phenomenon read at the level of the pooled estimate where country slopes range from -0.240 in Poland to 0.156 in Portugal, and dispersion within each group exceeds the gap between groups, so an average across sixteen countries averages responses too heterogeneous for the average to carry information. The CCEMG estimator allows country-specific loadings on unobserved common factors and is therefore the appropriate summary of a panel spanning two monetary regimes. It returns -0.076 (SE 0.046) on the extended sample, distinguishable from zero at the 10% level two-sided, and at the 5% level under the directional alternative of H1. This estimator remains valid when the unobserved common factors are themselves non-stationary (Kapetanios et al., 2011), so that cross-sectional dependence and factor non-stationarity are addressed jointly rather than in sequence. The power interpretation of Section 3.3 predicted this, though the gain combines added precision with a compositional shift toward the more negative non-eurozone slopes, and is therefore not a clean power gain on the eurozone estimand. One further result bears on the baseline design. The Pesaran CD test no longer rejects on the extended panel ($z = -1.29$, $p = 0.198$). Part of that attenuation is mechanical, the within-year benchmark falling from -0.100 to -0.067 as N rises, so the comparison is suggestive rather than decisive.

3.8. General discussion

Sections 3.3 to 3.7 established a single result many times over. Sovereign financial stress exhibits no detectable linear, lagged, moderated or non-linear effect on public cultural expenditure in this panel, and this null survives ten variations in measure, specification, sample and inference. Read narrowly, this is a series of failures to reject H_0. Read against the theoretical framework of Section 2.3, it is more informative than that, because it discriminates between two different pictures of how a financial crisis reaches a cultural budget.

The transmission mechanism set out in Section 2.3.1 is implicitly a gradient. A design built on a continuous stress index is well suited to detect exactly this kind of proportional response, and it does not find one. What the design does find, concentrated in Sections 3.6 and 3.7, is variation located in a small number of country-years rather than spread across the range of stress. The concave point estimate on the quadratic term, insignificant at the Wald test (p = 0.512) and bootstrap (p = 0.618), and sign-reversing when Greece is excluded, the leave-one-country-out reversal, and the sample-restriction result reported alongside R1 all point to the same handful of EU-IMF programme observations. These results generate a different, more political account of the same mechanism, they do not test it. Testing it would require an event-study design anchored on programme entry, at quarterly or budget-execution frequency. A rupture is an abrupt reallocation that an annual panel averaging over nineteen waves smooths away, and only execution-frequency data around the entry date can resolve it. Governments facing rising but still-manageable stress can defer adjustment. Spending-based consolidation is politically costly whatever its size (Ponticelli & Voth, 2020), and cultural cuts, however administratively easy, buy little fiscal space relative to the unrest they provoke; the incentive is to postpone, not to trim continuously. What forces an abrupt reallocation is not stress as such but the loss of market access, the condition that defined the programme countries, and one a continuous index registers only as an extreme value in its own tail. Bonet and Donato's (2011) reading of 2008 as a structural rupture rather than a cyclical episode fits this better than a smooth dose-response model. Most of the panel, on this account, experienced stress without rupture; the null for H1, H2 and H4 is exactly what a rupture-based mechanism produces when the ruptures are few and concentrated.

Figure 2
Country-specific slopes of cultural expenditure on SovCISS, 2006-2024

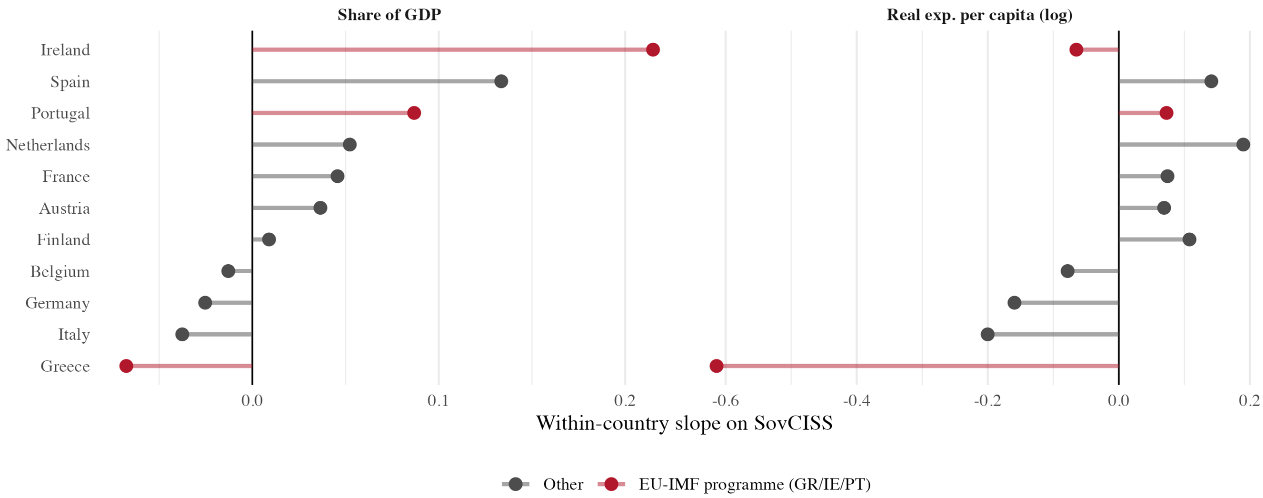

Figure 2 decomposes the association country by country on both measures. The distribution it reveals is the difficulty, not the evidence. Of the three programme states, only Greece is negative on both. Its nominal cultural budget fell from €388m in 2008 to €177m in 2010, a 54% contraction deep enough to drive the ratio down against a shrinking GDP, and its real per-capita slope is three times the next most negative in the panel. Ireland's GDP-share slope is positive and turns negative once the ratio is replaced by real expenditure per capita. Portugal's remains positive on both while three states that never entered a programme, Belgium, Germany and Italy, are also negative on both, at a third of the Greek magnitude or less. Loss of market access therefore does not by itself separate the countries that cut from those that did not. One programme state behaves as a threshold account predicts, one does so only once the denominator is cleaned and one does not at all. If a threshold exists, it lies somewhere beyond programme entry, at a depth only Greece reached, i.e. a claim about one country, not a regime this panel can identify. This is why the leave-one-country-out reversal of Section 3.7 is carried by Greece alone, not by the programme group.

This reframing carries a specific implication: built heritage does not depreciate smoothly under any circumstances. Deferred restoration compounds through the irreversibility mechanism established in Section 2.1.2, whether the deferral results from a proportional budget trim or from a single abrupt reallocation. A rupture-based transmission mechanism does not soften this concern relative to a gradient-based one, it relocates it. If the affected episodes are few, the panel's average null is compatible with heritage risk that is not diffusely distributed across all financially stressed states but concentrated precisely in the handful of country-years where external constraint left no room for the deferral strategy that other governments could afford. The finding that most of the sample shows no measurable response is not, on this reading, reassuring for heritage preservation. It says only that most of the sample was never asked to make the choice that damages heritage irreversibly.

The institutional-moderation result of Section 3.4 sits comfortably within this account. The marginal effect of stress is most negative at low institutional quality, attenuates toward the median, and turns marginally positive at high institutional quality, exactly the ordering the resilience literature predicts (Getzner, 2015; Rizzo & Throsby, 2006), without the panel being able to confirm it statistically. If resilience operates mainly by widening the menu of adjustment options available before a rupture becomes unavoidable, rather than by dampening a continuous marginal response to stress, then a panel of eleven countries observed annually may simply lack the resolution to detect it. The mechanism the case-study literature documents (Bonet & Donato, 2011; Getzner, 2015) operates within specific crisis episodes and specific institutional configurations, a form of heterogeneity an annual cross-country panel averages away rather than isolates. This is a reason to read the present null as compatible with the literature, in the same way that the MDE bounds established in Section 3.3 keep the null from being read as evidence that institutions do not matter.

On the rupture account, the encompassing result of Section 3.7 restates the Section 2.3.4 case from the opposite direction rather than contradicting it. If the mechanism runs through a few discrete episodes, binary indicators for those episodes and the extreme tail of the continuous index capture much the same variation, which is exactly what Section 3.6 showed in finding curvature and a programme-country effect inseparable. That overlap is what the threshold reading predicts, not independent evidence for it, and does not discriminate between it and the alternatives of Chapter 4.

Taken as a whole, the empirical chapter does not support the proposition that European states cut tangible heritage protection in proportion to the financial stress they face. It supports a narrower and more specific one. Measurable disruption to the public expenditure protecting tangible heritage is concentrated in the small set of episodes where financial stress left governments no room to manage the cut politically. Outside those episodes, cultural expenditure moves too slowly and too independently of annual financial conditions for a continuous stress gradient to leave a visible trace. The following chapter examines how far this conclusion can be pushed given the limitations of the design that produced it.

4. Limitations

The empirical strategy tests the effect of financial stress on cultural expenditure at the COFOG level of "cultural services" (sub-function 08.2), not on heritage-specific expenditure as such. No cross-country annual series isolates the conservation of tangible heritage from the broader cultural budget it sits within. The closest disaggregation, the ESSnet-Culture programme, is too fragmentary to support panel estimation. The thesis bridges this gap theoretically, through the cost-disease and irreversibility arguments of Section 2.1.2 that apply to built heritage specifically, rather than empirically. The results therefore speak directly to the public expenditure environment within which heritage preservation is funded, and only through that theoretical bridge to conservation itself. A finding of no aggregate effect on cultural-services expenditure does not rule out a more concentrated effect on the restoration and maintenance sub-component, the margin the third abandoned hypothesis of Section 3.2 was designed, and found underpowered, to test.

Two further boundaries follow from data availability rather than design. The restriction to eleven eurozone countries excludes states for which SovCISS exists but the transmission mechanism under test, operating through a shared monetary framework without independent exchange-rate adjustment, does not apply in the same way. Thus, the findings should not be extended to non-eurozone Europe without that caveat. Within the sample, the country-level CISS used as a robustness check is unavailable for Greece, precisely the country whose exclusion is shown through two independent routes in Section 3.7 to drive most of the instability in the baseline estimate. So the check that would speak most directly to Greece's role cannot include Greece, and SovCISS remains the only bridge across that gap.

A measurement property of the central explanatory variable, documented in Section 3.1, constrains interpretation more than estimation. The probability integral transform applied to the SovCISS components places each country's stress history on a common unit interval, so that the index is ordinally rather than cardinally comparable across states. It establishes that a country-year was unusual by that country's standards, not that it was more severe than a country-year recording a lower value elsewhere. Country fixed effects absorb the level of each national distribution, which leaves within-country identification intact and is part of why the design was built around them. What the property does bound is the cross-country content of β_1, estimated from differences in relative rather than absolute exposure, and it plausibly contributes to the attenuation of that coefficient. A specification asked to detect an effect of severity, using a regressor from which cross-country severity differences have been removed by construction, works against its own measure. No alternative was available within this sample. Raw sovereign spreads are cardinally comparable but capture one of the six stress symptoms the SovCISS integrates, and no harmonised cardinal composite exists for the eleven countries over this window.

The moderating role of institutional quality (Section 3.4) is measured through the World Bank's Government Effectiveness indicator rather than through decentralisation directly, the mechanism the theoretical literature actually emphasises (Van der Ploeg, 2006; Rizzo & Throsby, 2006). This substitution was not a free choice. A decentralisation index would have to be built from the central/local split of COFOG expenditure at the level of cultural services specifically, annually, and for all eleven countries. The problem is that combination of functional detail, frequency and coverage is not available, and the aggregate splits that do exist for the wider recreation, culture and religion division cannot be disaggregated to the sub-function tested here. A fiscal decentralisation measure is in any case close to time-invariant over a nineteen-year window, leaving it collinear with the country fixed effects alongside which it would have to be estimated, and it is itself plausibly responsive to the same fiscal pressures whose effect it is meant to moderate. WGI-GE is a defensible proxy for institutional capacity, but it is not decentralisation, and the resilience mechanism specifically attributed to multi-level government in Section 2.3.3 is accordingly tested only at one remove.

Persistence is tested in Section 3.5 through directly observed lags of the stress variable rather than through a dynamic specification with the lagged dependent variable, the form in which the entrenchment mechanism of Almeda et al. (2018) and Jaén-García (2021) is usually theorised: budgets that stay depressed after stress recedes, rather than expenditure that responds to stress with a delay. A dynamic panel would estimate that mechanism directly, but is exposed to Nickell (1981) bias at only nineteen annual waves per country. This bias does not shrink with the cross-sectional dimension and the standard corrections, Arellano-Bond and Blundell-Bond, are themselves ill-suited to resolve with only eleven cross-sectional units (Arellano & Bond, 1991; Blundell & Bond, 1998). The lagged-stress specification is the more conservative choice under this constraint, but it means the entrenchment mechanism is inferred from the persistence of the expenditure series and the absence of a lagged stress effect, rather than estimated directly. The bias at issue is specific to the dynamic form, demeaning correlates the lagged dependent variable with the transformed error, and the static baseline, carrying no such term, does not inherit it. The small-sample bias documented since Nickell (1981) therefore bears on the dynamic panel declined above, not on the specification actually estimated, as the simulation that follows makes precise.

A Monte Carlo calibrated to the panel (N = 11, T = 19, error autocorrelation 0.795) separates two concerns that the objection to small-N inference tends to conflate: whether the coefficient is biased, and whether inference around it is correctly sized. The first turns on strict exogeneity, the second on the number of clusters. Under strict exogeneity the two-way FE estimator is essentially unbiased (bias = 0.012), the within transformation eliminating the fixed effects algebraically whatever the number of countries, and country-clustered inference is correctly sized (rejection rate 0.059 against a nominal 0.05). This suggests that eleven clusters are not in themselves disqualifying. Once the regressor is allowed to respond to past shocks both properties fail, the bias rises to 0.142 and the rejection rate to 0.52. That figure comes from a feedback parameter chosen to stress the design rather than estimated here, and is a conservative bound rather than the bias of this panel. It does not diminish as the cross-section grows, so additional countries would not remove it, and it is of the same order as the minimum detectable effect of 0.17 reported in Section 3.3. The nulls are therefore consistent with a genuine absence of association and with a real effect offset by a bias of comparable size. The lead test of Section 3.7 rejects strict exogeneity at the eleven-country level, so the reading in which a bias offsets a real effect cannot be excluded on this basis alone. The violation is established, its empirical magnitude appears modest.

Two limitations established elsewhere in the analysis are noted here without being re-argued. At the 10% level, the encompassing exercise of Section 3.7 found discrete crisis dummies jointly significant net of SovCISS, which qualifies without overturning the case made in Section 2.3.4 for a continuous measure. And the two-way clustered standard errors of Section 3.7 rest on only nineteen year clusters, Driscoll-Kraay rests on the same time-series asymptotics. Neither is read as primary, the enumerated bootstrap, which requires no such asymptotics, is.

Beyond these design-level compromises, the identification strategy remains that of an observational panel. Country and year fixed effects absorb time-invariant heterogeneity and shocks common to the monetary union. None of them addresses the general possibility of a time-varying confounder correlated with both sovereign stress and the political willingness to protect cultural budgets, for which no instrument is available in this setting.

Finally, the COFOG series used throughout aggregates spending across all levels of government within each country. As noted in Section 2.3.2, this can mask offsetting dynamics, a central-government cut partly absorbed by subnational stability or the reverse, a concern sharpened by the asymmetric timing of central and local adjustment documented by Srakar et al. (2017) and Čopič et al. (2013). A government-level breakdown was not pursued, both for consistency with the annual frequency the panel requires and because it is unavailable for several sample countries; the null results of Chapter 3 should accordingly be read as describing consolidated general-government cultural expenditure, not the behaviour of any single tier within it.

5. Conclusion

This thesis has examined how systemic financial stress affects the public expenditure through which European states preserve their tangible cultural heritage. The motivating argument, developed in Chapter 2, is that heritage occupies a particular position among the objects of cultural policy. Conservation resists productivity gains, since the regulatory protection of a listed structure forecloses the cheaper methods other sectors use to absorb budgetary pressure, so that a postponed restoration is not a neutral deferral but a compounding liability. Public budgets remain the dominant source of funding for that activity, and discretionary items are the first to be reduced when consolidation arrives. The empirical chapter tested whether the sequence actually runs through, using a balanced panel of eleven eurozone countries between 2006 and 2024 and a continuous, market-determined index of sovereign stress in place of the crisis dummies the literature normally employs. Five hypotheses, frozen before estimation, addressed the direct effect on two expenditure measures, moderation by institutional quality, persistence, and non-linearity.

None was confirmed on the specification against which it was frozen before estimation. The finding is not that cultural budgets were unaffected, sovereign stress does track the budgetary share of culture when it is allowed to operate through the channels it actually runs through. The association is negative and significant before growth, debt and unemployment enter and indistinguishable from zero once they do. What the panel does not find is a compositional response net of the macro-fiscal cycle, or a response proportional to stress across the observed range. Culture fell with the rest of the budget but was not singled out; and where it was singled out, the constraint had already become absolute. What the panel isolates, instead of a proportional response, is a discontinuity confined to the few country-years in which a state lost access to market financing, variation that Greece carries almost single-handedly. My own reading is that the discrete-crisis literature, and still more the policy commentary drawing on it, has generalised from a small number of programme countries to a continental phenomenon the aggregate data do not support.

On a narrow criterion of fit, the substitution from discrete dummies to a continuous measure gained little. I would nonetheless defend it, on the ground that fit was never the argument. A dummy can establish that budgets fell after a date the researcher selected, whereas only a continuous index can ask whether they fall in proportion to stress, above a threshold, or with a lag. The null answers obtained here are answers to questions a categorical treatment of the shock is not equipped to pose, and the encompassing result itself becomes interpretable only because both representations were estimated within a single specification.

The second contribution concerns inference, and here I would make a stronger claim. Cross-country panels in cultural economics commonly rest on ten to twenty clusters, in which the reliability of asymptotic inference has to be verified rather than assumed, and they rarely report non-rejections with any indication of what magnitude of effect the design could have detected. Every null established here is paired with an exact wild cluster bootstrap computed by full enumeration of all 2048 sign vectors, and with a minimum detectable effect expressed as a share of the average cultural budget. The design can only detect stress-induced contractions above roughly a quarter, depending on the normalisation, of cultural spending over an episode of typical amplitude, and excludes nothing below it. A null reported without such bounds is not a finding but a silence, and the field's tolerance for the latter has allowed underpowered designs and genuine absences of effect to be presented in identical terms.

The third contribution concerns a measurement practice extending well beyond this thesis. Cultural expenditure normalised by GDP has a procyclical denominator, and Section 3.8 established that this is not a technicality. The consequence is visible twice: the sign of the stress response reverses for one programme country once the ratio is replaced by real expenditure per capita, and the per-capita specification is the only one of the five specifications re-estimated on the extended panel that preserves its sign when the panel is extended. I would therefore argue that the share of GDP should not be read as a measure of cultural policy effort in recessionary periods without a level check in real per-capita terms. Comparative statements built on the ratio during a downturn describe the denominator at least as much as the policy, and this holds irrespective of whether the hypotheses tested here are accepted.

That most of the panel shows no measurable response is not, on reflection, reassuring. Deferred maintenance compounds irreversibly whether the deferral follows a proportional trim or an abrupt reallocation, and the null establishes only that most sample countries were never forced into the second. Two proposals follow. The first is institutional, if the risk to heritage is concentrated in ruptures rather than distributed along a gradient, protective arrangements must bind before discretion disappears, since by the time a state loses market access its capacity for domestic protection has closed with it. Multiannual conservation commitments, envelopes ring-fenced against in-year reallocation, and co-financing organised supranationally rather than nationally all have that property. I cannot demonstrate that such instruments work, and the case for them rests on where the risk sits rather than on evidence of their effect. The second proposal is statistical, and I regard it as the more pressing of the two. No harmonised European series separates the conservation of built heritage from cultural services as a whole, which is why the link between the expenditure studied here and preservation itself remains a theoretical bridge rather than an estimated one. A disaggregation of COFOG 08.2 comparable across member states would do more for this question than any refinement of estimator.

Two extensions follow from Chapter 4. The first is the event-study design set out in Section 3.8, which the rupture reading makes the natural next step. The second addresses that expenditure is only a proxy for preservation, pairing budget series with the condition data some national heritage-at-risk registers already compile would let the irreversibility argument motivating this thesis be tested rather than assumed.

The question was whether financial crises disturb the intergenerational transmission of the built past. Most European states, most of the time, did not visibly make that trade-off under financial pressure; those that did were left without an alternative. Neither half of that sentence is reassuring alone. Together, they suggest that what protects heritage in a crisis is less the resilience of cultural policy than the fiscal capacity that keeps the choice from ever being forced.

Bibliography

Adie, B. A., Falk, M., & Savioli, M. (2020). Overtourism as a perceived threat to cultural heritage in Europe. *Current Issues in Tourism*, *23*(14), 1737-1741. https://doi.org/10.1080/13683500.2019.1687661

Albertelli, A., & Mignosa, A. (2026). Public-private partnerships: When governance tools stimulate alternative funding of the heritage sector and beyond. In C. Dalla Chiesa & A. Rykkja (Eds.), *Cultural Funding and Financing. Cultural Economics & the Creative Economy*, 321-344. Palgrave Macmillan. https://doi.org/10.1007/978-3-031-96696-5_14

Alesina, A., Favero, C., & Giavazzi, F. (2019). Effects of austerity : Expenditure- and tax-based approaches. *Journal of Economic Perspectives*, *33*(2), 141-162. https://doi.org/10.1257/jep.33.2.141

Almeda, P., Sagarra, A., & Tataret, M. (2018). *Public spending on culture in Europe 2007-2015*. European Parliament. https://www.europarl.europa.eu/cmsdata/139228/ALMEDA_Public%20spending%20on%20culture%202007-2015%20-%2021F%20-%20EP%20-Cuture%20Commission%2021%20FEB.pdf

Anderson, T. W., & Hsiao, C. (1981). Estimation of dynamic models with error components. *Journal of the American Statistical Association, 76*(375), 598–606. https://doi.org/10.1080/01621459.1981.10477691

Arellano, M., & Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *The Review of Economic Studies, 58(*2), 277–297. https://doi.org/10.2307/2297968

Baumol, W. J., & Bowen, W. G. (1966). Performing arts: The economic dilemma. A study of problems common to theater, opera, music, and dance. Twentieth Century Fund. https://doi.org/10.1017/S0770451800055287

Benhamou, F. (1996). Is increased public spending for the preservation of historic monuments inevitable? The French case. *Journal of Cultural Economics, 20(2)*, 115-131. https://doi.org/10.1007/s10824-005-5131-y

Benhamou, F. (2013). Public intervention for cultural heritage: normative issues and tools. In I. Rizzo & A. Mignosa (Eds.), *Handbook on the Economics of Cultural Heritage* (pp. 3-16). Edward Elgar Publishing. https://doi.org/10.4337/9780857931009.00008

Benhamou, F. (2019). *Économie du patrimoine culturel*. La Découverte. https://doi.org/10.3917/dec.benha.2019.01

Benhamou, F., & Thesmar, D. (2011). *Valoriser le patrimoine culturel de la France* (Rapport CAE n°97). La Documentation française. https://cae-eco.fr/static/pdf/097.pdf

Bergé, L. (2018). *Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm*. (DEM discussion paper No. 2018-13) hal-05654564

Bertacchini, E., & Dalle Nogare, C. (2014). Public provision vs. outsourcing of cultural services : Evidence from Italian cities. *European Journal of Political Economy*, *35*, 168-182. https://doi.org/10.1016/j.ejpoleco.2014.06.004

Betzler, D., Loots, E., Prokůpek, M., Marques, L., & Grafenauer, P. (2020). COVID-19 and the arts and cultural sectors : Investigating countries' contextual factors and early policy measures. *International Journal of Cultural Policy*, *27*(6), 796-814. https://doi.org/10.1080/10286632.2020.1842383

Bloom, H. S. (1995). Minimum detectable effects: A simple way to report the statistical power of experimental designs. *Evaluation Review, 19*(5), 547–556. https://doi.org/10.1177/0193841X9501900504

Blundell, R., & Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics, 87*(1), 115–143. https://doi.org/10.1016/S0304-4076(98)00009-8

Bonet, L. (2026). Direct government contribution to culture: Legitimacy, financial mechanisms and level of spending compared internationally. In C. Dalla Chiesa & A. Rykkja (eds) *Cultural Funding and Financing*. *Cultural Economics & the Creative Economy*. (pp. 35-64). Palgrave Macmillan, Cham. https://doi.org/10.1007/978-3-031-96696-5_2

Bonet, L., & Donato, F. (2011). The financial crisis and its impact on the current models of governance and management of the cultural sector in Europe. *European Journal of Cultural Management and Policy*, *1*. https://doi.org/10.3389/ejcmp.2023.v1iss1-article-1

Breunig, C., & Busemeyer, M. R. (2012). Fiscal austerity and the trade-off between public investment and social spending. *Journal of European Public Policy*, *19*(6), 921-938. https://doi.org/10.1080/13501763.2011.614158

Buchanan, J. M. (1965). An economic theory of clubs. *Economica*, *32*(125), 1-14. https://doi.org/10.2307/2552442

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors, *The Review of Economics and Statistics, 90*(3), 414-427. https://doi.org/10.1162/rest.90.3.414

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering, *Journal of Business & Economic Statistics, 29*(2), 238-249 https://doi.org/10.1198/jbes.2010.07136

Cominelli, F., & Greffe, X. (2019). L'économie politique du patrimoine culturel. *In Situ. Au regard des sciences sociales*, *1*. https://doi.org/10.4000/insituarss.436

Čopič, V., Inkei, P., Kangas, A., & Srakar, A. (2013). *Trends in public funding for culture in the EU*. (EENC Report, revised July 2014). European Expert Network on Culture.

Council of Europe (2003), *Tax incentives for cultural heritage conservation* (Report No. 9913). Council of Europe Parliamentary Assembly: Committee on Culture, Science and Education. https://assembly.coe.int/nw/xml/XRef/Xref-XML2HTML-en.asp?fileid=17166&lang=en

Cuccia, T. (2012). Is it worth being inscribed in the World Heritage List? A case study of 'The Baroque Cities in Val di Noto' (Sicily). *Rivista Italiana di Economia Demografia e Statistica, 66*(2), 169-190. https://www.sieds.it/listing/RePEc/journl/2012LXVI_N2_12.pdf

Dalle Nogare, C., & Galizzi, M. M. (2011). The political economy of cultural spending : Evidence from Italian cities. *Journal of Cultural Economics*, *35*(3), 203-231. https://doi.org/10.1007/s10824-011-9145-3

De Wit, A., & Bekkers, R. (2016). Government support and charitable donations: A meta-analysis of the crowding-out hypothesis. *Journal of Public Administration Research and Theory, 27*(2), 301-319. https://doi.org/10.1093/jopart/muw044

De Wit, A., Bekkers, R., & Broese Van Groenou, M. (2016). Heterogeneity in crowding-out : When are charitable donations responsive to government support? *European Sociological Review, 33*(1), 59-71. https://doi.org/10.1093/esr/jcw048

Di Pietro, F., & Butticè, V. (2020). Institutional characteristics and the development of crowdfunding across countries. *International Review of Financial Analysis, 71*, Article 101543. https://doi.org/10.1016/j.irfa.2020.101543

Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics, 80*(4), 549–560. https://doi.org/10.1162/003465398557825

ESPON. (2020). *Measuring the economic impact of cultural heritage at territorial level*. ESPON. https://archive.espon.eu/sites/default/files/attachments/Working%20Paper,%20Cultural%20heritage.pdf

Freitag, M., & Vatter, A. (2008). Decentralization and fiscal discipline in sub-national governments: Evidence from the Swiss Federal System. *Publius: The Journal of Federalism, 38*(2), 272-294. https://doi.org/10.1093/publius/pjm038

Garcia-de-Andoain, C., & Kremer, M. (2018). *Beyond spreads: Measuring sovereign market stress in the euro area* (ECB Working Paper No. 2185). European Central Bank. https://doi.org/10.2866/329712

Getzner, M. (2002). Determinants of public cultural expenditures : An exploratory time series analysis for Austria. *Journal of Cultural Economics, 26*(4), 287-306. https://doi.org/10.1023/A:1019976717566

Getzner, M. (2015). Cultural politics : Exploring determinants of cultural expenditure. *Poetics*, *49*, 60-75. https://doi.org/10.1016/j.poetic.2015.02.001

Gootjes, B., & De Haan, J. (2020). Procyclicality of fiscal policy in European Union countries. *Journal of International Money and Finance*, *120*, 102276. https://doi.org/10.1016/j.jimonfin.2020.102276

Greffe, X. (2010). Introduction : L'économie de la culture est-elle particulière ? *Revue d'économie politique, 120*(1), 1-34. https://doi.org/10.3917/redp.201.0001

Greffe, X. (2013). Existe-t-il une économie du patrimoine culturel ? *Revue française de finances publiques*, *N° 122*(2), 23-35. https://doi.org/10.3917/rffp.122.0023

Holló, D., Kremer, M., & Lo Duca, M. (2012). *CISS — A composite indicator of systemic stress in the financial system* (ECB Working Paper No. 1426). European Central Bank. https://doi.org/10.2139/ssrn.2018792

Husson, J.-F. (2013). Une approche sociopolitique du financement de la culture. In H. Pauliat (Ed.), *Culture et politiques publiques culturelles en Europe : quelles valeurs à préserver en temps* *de crise ?* (pp. 111-126). Presses universitaires de Limoges. https://hdl.handle.net/2268/134107

Jaén-García, M. (2021). Displacement effect and ratchet effect: Testing of two alternative hypotheses. *Sage Open*, *11*(1), 21582440211003577. https://doi.org/10.1177/21582440211003577

Jelinčić, D. A., & Šveb, M. (2021). Financial sustainability of cultural heritage : A review of crowdfunding in Europe. *Journal of Risk and Financial Management*, *14*(3), 101. https://doi.org/10.3390/jrfm14030101

Kapetanios, G. (2003). *Determining the poolability of individual series in panel datasets* (Working Paper No. 499). Queen Mary, University of London, Department of Economics. https://www.qmul.ac.uk/sef/media/econ/research/workingpapers/2003/items/wp499.pdf.

Kapetanios, G., Pesaran, M. H., & Yamagata, T. (2011). Panels with non-stationary multifactor error structures. *Journal of Econometrics, 160*(2), 326–348. https://doi.org/10.1016/j.jeconom.2010.10.001

Kaufmann, D., Kraay, A., & Mastruzzi, M. (2011). The Worldwide Governance Indicators: Methodology and analytical issues. *Hague Journal on the Rule of Law, 3*(2), 220–246. https://www.researchgate.net/publication/215645597_The_Worldwide_Governance_Indicators_Methodology_and_Analytical_Issues

Kostakis, I., & Lolos, S. (2024). Uncovering the impact of cultural heritage on economic growth : Empirical evidence from Greek regions, 2000-2019. *The Annals of Regional Science*, *73*(3), 1209-1239. https://doi.org/10.1007/s00168-024-01280-3

Lundberg, J. (2006). Spatial interaction model of spillovers from locally provided public services. *Regional Studies, 40*(6), 631–644. https://doi.org/10.1080/00343400600868788

Macdonald, S., & Cheong, C. (2014). *The role of public-private partnerships and the third sector in conserving heritage buildings, sites, and historic urban areas*. The Getty Conservation Institute. *https://www.getty.edu/publications-reports/item/2K7M3B*

MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, *232*(2), 272-299. https://doi.org/10.1016/j.jeconom.2022.04.001

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica, 49*(6), 1417–1426. https://doi.org/10.2307/1911408

OECD. (2022). *The culture fix : Creative people, places and industries*. OECD Publishing. https://doi.org/10.1787/991bb520-en

Osborne, S. P. (2006). The New Public Governance?. *Public Management Review*, *8*(3), 377-387. https://doi.org/10.1080/14719030600853022

Peacock, A. T., & Wiseman, J. (1961). The growth of public expenditure in the United Kingdom. Princeton University Press. http://www.nber.org/books/peac61-1

Pesaran, M. H. (2004). General diagnostic tests for cross-sectional dependence in panels. (CESifo Working Paper No. 1229). *CESifo*. http://dx.doi.org/10.2139/ssrn.572504

Pesaran, M. H. & Smith, R. (1995). Estimating long run relationships from dynamic heterogeneous panels. *Journal of Econometrics, 68*(1), 79-113. https://doi.org/10.1016/0304-4076(94)01644-F

Ponticelli, J., & Voth, H.-J. (2020). Austerity and anarchy : Budget cuts and social unrest in Europe, 1919-2008. *Journal of Comparative Economics*, *48*(1), 1-19. https://doi.org/10.1016/j.jce.2019.09.007

Potrafke, N. (2011). Public expenditures on education and cultural affairs in the west german states : Does government ideology influence the budget composition? *German Economic Review*, *12*(1), 124-145. https://doi.org/10.1111/j.1468-0475.2010.00507.x

Pratt, A. C. (2017). Beyond resilience : Learning from the cultural economy. *European Planning Studies*, *25*(1), 127-139. https://doi.org/10.1080/09654313.2016.1272549

Reinhart, C. M., & Rogoff, K. S. (2011). From Financial crash to debt crisis. *American Economic Review*, *101*(5), 1676-1706. https://doi.org/10.1257/aer.101.5.1676

Rizzo, I., & Throsby, D. (2006). Economic analysis and Public Policy. In V. Ginsburg & D. Throsby (Eds). *Handbook of the Economics of Art and Culture* (pp. 983-1016). Elsevier. https://doi.org/10.1016/S1574-0676(06)01028-3

Rubio-Arostegui, J. A., & Rius-Ulldemolins, J. (2020). Cultural policies in the South of Europe after the global economic crisis : Is there a Southern model within the framework of

European convergence? *International Journal of Cultural Policy*, *26*(1), 16-30. https://doi.org/10.1080/10286632.2018.1429421

Samuelson, P. A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics*, *36*(4), 387. https://doi.org/10.2307/1925895

Seaman, B. A. (2013). The role of the private sector in cultural heritage. In I. Rizzo & A. Mignosa (Eds.), *Handbook on the Economics of Cultural Heritage* (pp. 111-125). Edward Elgar Publishing. https://doi.org/10.4337/9780857931009.00013

Škrabić Perić, B., Šimundić, B., Muštra, V., & Vugdelija, M. (2021). The role of UNESCO cultural heritage and cultural sector in tourism development : The case of EU countries. *Sustainability*, *13*(10), 5473. https://doi.org/10.3390/su13105473

Srakar, A., Vecco, M., & Tóth, Á. (2017). The tale of the cuts and raises: public budgets for culture in the European countries during the financial crisis. *Revista Hacienda Pública Española*, *221*(2), 83-109. https://www.ief.es/docs/destacados/publicaciones/revistas/hpe/221_Art4.pdf

Throsby, D. (2001). *Economics and culture*. Cambridge University Press. https://doi.org/10.1017/CBO9781107590106

Throsby, D. (2010). *The Economics of cultural policy*. Cambridge University Press. https://doi.org/10.1017/CBO9780511845253

UNESCO. (1972). *Convention concerning the protection of the world cultural and natural heritage*. https://whc.unesco.org/en/conventiontext/

Van der Ploeg, F. (2006). The making of cultural policy: A European perspective. In V. Ginsburg & D. Throsby (eds.). *Handbook of the Economics of Art and Culture* (pp. 1183-1221). Elsevier. https://doi.org/10.1016/S1574-0676(06)01034-9

Vecco, M. (2010). A definition of cultural heritage : From the tangible to the intangible. *Journal of Cultural Heritage*, *11*(3), 321-324. https://doi.org/10.1016/j.culher.2010.01.006

Vincent, A., & Wunderle, M. (2010). Le financement privé de la culture: *Dossiers du CRISP*, *74*(1), 9-84. https://doi.org/10.3917/dscrisp.074.0009

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT Press. https://ipcid.org/evaluation/apoio/Wooldridge%20-%20Cross-section%20and%20Panel%20Data.pdf

Appendices

Appendix

Table 1 - Descriptive statistics

Variable	N	Mean	SD	Min	P25	Median	P75	Max
Cultural exp. (% GDP)	209	0.425	0.153	0.079	0.317	0.445	0.530	0.717
Cultural exp. (% total gov. exp.)	209	0.874	0.277	0.149	0.735	0.918	1.046	1.667
SovCISS (annual mean)	209	0.218	0.211	0.029	0.059	0.125	0.317	0.946
Real GDP growth (%)	209	1.329	3.943	-10.900	0.100	1.600	2.800	24.600
Public debt (% GDP)	209	92.255	36.305	23.700	65.400	86.300	111.200	209.400
Unemployment rate (%)	209	9.038	5.055	3.100	5.600	7.800	10.300	27.500

Notes:

Balanced panel: 11 eurozone countries, 2006-2024 (N = 209). Sources: Eurostat COFOG (cultural expenditure), ECB SDW (SovCISS), Eurostat (macroeconomic controls).

Table 2a - Pairwise correlations between time-varying regressors

Variable	SovCISS (mean)	GDP growth	Public debt (% GDP)	Unemployment rate
SovCISS (mean)	1.000	-0.368	0.244	0.431
GDP growth	-0.368	1.000	-0.189	-0.183
Public debt (% GDP)	0.244	-0.189	1.000	0.593
Unemployment rate	0.431	-0.183	0.593	1.000

Notes:

Pearson correlations. N = 209.

Table 2b - Variance inflation factors

Variable	VIF
SovCISS (mean)	1.38
GDP growth	1.17
Public debt (% GDP)	1.56
Unemployment rate	1.79

Notes:

VIF computed on a pooled OLS regression of cult_exp_gdp on all time-varying regressors. Conventional threshold of concern: VIF > 5 or > 10.

Table 3a - Baseline: financial stress and cultural expenditure

	H1 - M1	H1 - M2	H1 - M3	H1 - DK(2)	H2 - M1	H2 - M2	H2 - M3	H2 - DK(2)
SovCISS (mean)	−0.013	−0.024	−0.025	−0.025	−0.123**	−0.105*	−0.036	−0.036
	(0.035)	(0.029)	(0.062)	(0.024)	(0.048)	(0.051)	(0.110)	(0.066)
GDP growth		−0.002	−0.002	−0.002		0.002	0.003	0.003
		(0.001)	(0.001)	(0.002)		(0.002)	(0.002)	(0.002)
Public debt (% GDP)			0.000	0.000			−0.001	−0.001
			(0.001)	(0.001)			(0.001)	(0.001)
Unemployment rate			−0.000	−0.000			−0.004	−0.004
			(0.004)	(0.004)			(0.009)	(0.006)
Num.Obs.	209	209	209	209	209	209	209	209
R2	0.928	0.929	0.929	0.929	0.910	0.911	0.918	0.918
R2 Within	0.002	0.009	0.014	0.014	0.039	0.043	0.117	0.117

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs unless noted. Country and year fixed effects in all models.

Table 3b - Wild cluster bootstrap: enumerated p-values (H1-H5 and Table 4c)

Hypothesis	Coefficient	H_A (one-sided)	t (observed)	p (two-sided)	p (one-sided)
H1	β_1 (SovCISS)	$\beta_1 < 0$	-0.404	0.769	0.385
H2	β_1 (SovCISS)	$\beta_1 < 0$	-0.333	0.728	0.364
H3	γ_3 (SovCISS × WGI-GE), Table 4a M3	$\gamma_3 > 0$	1.113	0.462	0.231
H3	γ_3 (SovCISS × WGI-GE), Table 4c + programme	$\gamma_3 > 0$	1.757	0.105	0.052
H3	SovCISS × EU-IMF programme, Table 4c	two-sided	1.331	0.329	—
H4	γ_4 (SovCISS, t-1)	$\gamma_4 < 0$	-0.572	0.692	0.346
H5	γ_5 (SovCISS², centered)	two-sided	-0.681	0.618	—

Notes:

Exact enumeration over the 2,048 Rademacher sign vectors. H1, H2 and H4 are tested one-sided against a negative alternative, H3 is one-sided against a positive one. H5 is two-sided by construction.

Table 4a - H3: WGI-GE as moderator of financial stress

	Interaction only	+ WGI main	M3 baseline	M3 DK(2)
SovCISS (centered)	0.009	−0.017	−0.017	−0.017
	(0.042)	(0.032)	(0.047)	(0.042)
SovCISS x WGI-GE	0.048	0.042	0.045	0.045
	(0.062)	(0.060)	(0.041)	(0.056)
WGI-GE (centered)		−0.070	−0.070	−0.070***
		(0.042)	(0.043)	(0.023)
GDP growth			−0.002	−0.002
			(0.001)	(0.001)
Public debt (% GDP)			0.000	0.000
			(0.001)	(0.001)
Unemployment rate			−0.001	−0.001
			(0.004)	(0.003)
Num.Obs.	209	209	209	209
R2	0.929	0.933	0.933	0.933
R2 Within	0.013	0.061	0.072	0.072

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs unless noted. Country and year fixed effects.

Table 4b - H3: marginal effect of financial stress by institutional quality

Percentile of WGI-GE	WGI-GE	Marginal effect	SE	t	p-value
10th (weakest)	0.43	−0.057	0.067	−0.85	0.415
50th (median)	1.46	−0.011	0.046	−0.23	0.824
90th (strongest)	1.93	+0.011	0.047	+0.23	0.823

Marginal effect of SovCISS on cultural expenditure (% GDP), evaluated at the 10th, 50th and 90th percentiles of WGI-GE. Delta-method SEs from the country-clustered vcov of the H3 baseline (h3_full); two-sided p-values on t with G - 1 = 10 degrees of freedom. * p<0.1, ** p<0.05, *** p<0.01.

Table 4c - H3: control for EU-IMF programme interaction

	Baseline	+ SovCISS x EU-IMF programme
SovCISS (centered)	−0.017	−0.040
	(0.047)	(0.042)
WGI-GE (centered)	−0.070	−0.067
	(0.043)	(0.042)
GDP growth	−0.002	−0.001
	(0.001)	(0.001)
Public debt (% GDP)	0.000	0.000
	(0.001)	(0.001)
Unemployment rate	−0.001	−0.001
	(0.004)	(0.004)
SovCISS x WGI-GE	0.045	0.069
	(0.041)	(0.039)
SovCISS x EU-IMF programme		0.066
		(0.050)
Num.Obs.	209	209
R2	0.933	0.935
R2 Within	0.072	0.089

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs. Country and year fixed effects.

Table 4d - H3: country-level marginal effects implied by Table 4c

Country	WGI-GE (mean)	EU-IMF programme	Implied marginal effect	SE	p-value
Greece	0.36	Yes	−0.039	0.085	0.654
Italy	0.45	No	−0.100	0.059	0.124
Portugal	0.98	Yes	+0.003	0.073	0.967
Spain	1.03	No	−0.060	0.046	0.220
Ireland	1.41	Yes	+0.033	0.069	0.645
Belgium	1.44	No	−0.031	0.041	0.463
France	1.47	No	−0.030	0.041	0.484
Austria	1.67	No	−0.016	0.041	0.713
Germany	1.77	No	−0.009	0.042	0.835
Netherlands	1.87	No	−0.002	0.043	0.961
Finland	2.06	No	+0.011	0.046	0.817

Marginal effect of SovCISS on cultural expenditure (% GDP) implied by the coefficients of Table 4c, evaluated for each country at its own sample-period mean of WGI-GE. The programme group comprises the three sample countries placed under a full EU-IMF adjustment programme (Greece, Ireland, Portugal). Delta-method SEs from the country-clustered vcov of h3_troika; two-sided p-values on t with G - 1 = 10 degrees of freedom. Correlation between the programme indicator and centred WGI-GE: -0.447. * p<0.1, ** p<0.05, *** p<0.01.

Table 5 - H4: persistence of financial stress effects

	Contemporaneous	Lag only	Joint	Joint DK(2)	Two lags (n=187)
SovCISS (t)	−0.029		−0.017	−0.017	−0.023
	(0.060)		(0.059)	(0.026)	(0.055)
GDP growth	−0.001	−0.001	−0.001	−0.001	−0.001
	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)
Public debt (% GDP)	0.000	0.000	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Unemployment rate	0.001	0.001	0.001	0.001	0.004
	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
SovCISS (t-1)		−0.035	−0.026	−0.026	−0.041*
		(0.053)	(0.046)	(0.043)	(0.020)
SovCISS (t-2)					−0.010
					(0.053)
Num.Obs.	198	198	198	198	187
R2	0.929	0.929	0.929	0.929	0.937
R2 Within	0.015	0.016	0.018	0.018	0.044
Cumulative effect (2-year)			-0.0435		
95% CI			[-0.1815 ; 0.0945]		
Cumulative effect (3-year)					-0.0740
95% CI					[-0.2407 ; 0.0928]
Joint Wald (2-year, F)			F=0.23, p=0.800		
Joint Wald (3-year, F)					F=1.58, p=0.255

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs unless noted. Country and year fixed effects.

Table 6a - H5: non-linearity (quadratic and threshold spline)

	Linear	H5 : quadratic	DK(2)	Spline (q75)	Spline (q90)	H2 (% total exp)
SovCISS (centered)	−0.025	0.009	0.009			0.036
	(0.062)	(0.074)	(0.041)			(0.162)
GDP growth	−0.002	−0.002	−0.002	−0.002	−0.002	0.002
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.002)
Public debt (% GDP)	0.000	0.000	0.000	0.000	0.000	−0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Unemployment rate	−0.000	−0.001	−0.001	−0.001	−0.001	−0.005
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.009)
$SovCISS^2$ (centered)		−0.080	−0.080			−0.170
		(0.117)	(0.095)			(0.255)
SovCISS (raw)				0.010	0.002	
				(0.084)	(0.070)	
SovCISS - q75 (knot)				−0.051		
				(0.108)		
SovCISS - q90 (knot)					−0.082	
					(0.112)	
Num.Obs.	209	209	209	209	209	209
R2	0.929	0.930	0.930	0.929	0.930	0.918
R2 Within	0.014	0.020	0.020	0.017	0.020	0.123
AIC	−680.3	−679.6	−679.6	−678.8	−679.5	−400.3
BIC	−570.1	−565.9	−565.9	−565.2	−565.9	−286.7

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs unless noted. Country and year fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

Table 6b - H5: marginal effect of financial stress across the stress distribution

Percentile of SovCISS	SovCISS	Marginal effect	SE	t	p-value
10th	0.042	+0.037	0.103	+0.36	0.727
50th (median)	0.125	+0.024	0.088	+0.27	0.792
75th	0.317	−0.007	0.064	−0.10	0.919
90th	0.503	−0.036	0.066	−0.55	0.596

Marginal effect of SovCISS on cultural expenditure (% GDP), evaluated at the 10th, 50th, 75th and 90th percentiles of SovCISS. Delta-method SEs from the country-clustered vcov of the quadratic model; two-sided p-values on t with G - 1 = 10 degrees of freedom. * p<0.1, ** p<0.05, *** p<0.01.

Table 7a - Robustness battery (H1): financial stress and cultural expenditure

	Baseline	R1 CISS 10c	R2 CISS EA	R2 no-year-FE	R3 SovMax	R4 + Ideology	R5 log p.c.	R7 excl. COVID	R8 two-way cluster
Dependent variable	% GDP	% GDP	% GDP	% GDP	% GDP	% GDP	log real p.c.	% GDP	% GDP
SovCISS (mean)	−0.025			0.039		−0.022	−0.187	−0.026	−0.025
	(0.062)			(0.035)		(0.061)	(0.182)	(0.061)	(0.052)
GDP growth	−0.002	−0.003***	−0.001	−0.002*	−0.001	−0.001	0.005	−0.002	−0.002
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.004)	(0.001)	(0.001)
Public debt (% GDP)	0.000	0.000	−0.000	−0.001	0.000	0.000	−0.004	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Unemployment rate	−0.000	−0.003	0.001	−0.000	−0.001	−0.000	−0.007	−0.000	−0.000
	(0.004)	(0.002)	(0.002)	(0.001)	(0.003)	(0.003)	(0.009)	(0.004)	(0.004)
CISS (country)		−0.031							
		(0.059)							
CISS (euro area)			0.070**						
			(0.027)						
SovCISS (max)					−0.014				
					(0.039)				
Government ideology (ParlGov)						−0.009*			
						(0.005)			
log(real GDP p.c.)							−0.308		
							(0.200)		
Num.Obs.	209	190	209	209	209	207	209	187	209
R2	0.929	0.908	0.916	0.913	0.929	0.933	0.981	0.929	0.929
R2 Within	0.014	0.048	0.128	0.096	0.013	0.072	0.283	0.010	0.014
R6 jackknife range	[-0.0559; 0.0501]								

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs unless noted. Both R2 columns use country fixed effects only (no year FE). R8 uses two-way (country x year) clustering. The R5 coefficient is a semi-elasticity and is not comparable in magnitude with the other columns.

Table 7b - Robustness battery (H2): cultural expenditure (% total exp.)

	Baseline	R1 CISS 10c	R2 CISS EA	R3 SovMax	R4 + Ideo	R7 excl. COVID	R8 two-way cluster
SovCISS (mean)	−0.036				−0.030	−0.036	−0.036
	(0.110)				(0.113)	(0.109)	(0.099)
GDP growth	0.003	0.001	0.002*	0.002	0.004	0.002	0.003**
	(0.002)	(0.001)	(0.001)	(0.002)	(0.003)	(0.002)	(0.001)
Public debt (% GDP)	−0.001	−0.002	−0.003*	−0.001	−0.001	−0.002	−0.001
	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Unemployment rate	−0.004	−0.014***	−0.002	−0.004	−0.004	−0.004	−0.004
	(0.009)	(0.003)	(0.002)	(0.008)	(0.008)	(0.010)	(0.008)
CISS country		−0.095					
		(0.189)					
CISS eurozone			0.065*				
			(0.034)				
SovCISS (max)				−0.045			
				(0.074)			
Gov. ideology (ParlGov)					−0.017*		
					(0.008)		
Num.Obs.	209	190	209	209	207	187	209
R2	0.918	0.885	0.909	0.918	0.923	0.923	0.918
R2 Within	0.117	0.299	0.314	0.120	0.171	0.131	0.117
R6 jackknife range	[-0.1117; 0.1025]						

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs unless noted. R2 uses country FE only. R8 uses two-way clustering. R5 not applicable to the total-expenditure share.

Table 7c - Placebo: leads of SovCISS (strict-exogeneity check)

	One lead (t+1)	Two leads (t+1, t+2)
SovCISS (t)	−0.045	−0.023
	(0.034)	(0.048)
SovCISS (t+1)	0.023	−0.055***
	(0.055)	(0.014)
GDP growth	−0.002	−0.003*
	(0.001)	(0.001)
Public debt (% GDP)	0.000	0.000
	(0.001)	(0.001)
Unemployment rate	−0.000	−0.000
	(0.004)	(0.005)
SovCISS (t+2)		0.081
		(0.053)
Num.Obs.	198	187
R2	0.930	0.933
R2 Within	0.018	0.058
Joint Wald F (two leads)		13.85
p (analytic, F(2,10))		0.0013
p (enumerated bootstrap)		0.0010

* p < 0.1, ** p < 0.05, *** p < 0.01

Placebo test: future stress cannot cause current expenditure, so non-zero lead coefficients signal a violation of strict exogeneity. Country-clustered SEs; country and year fixed effects. The joint bootstrap resamples restricted residuals under the imposed null over all 2,048 Rademacher sign vectors. * p<0.1, ** p<0.05, *** p<0.01.

Table 8 - Discrete crisis dummies vs continuous stress measure (H1)

	Dummies (replication)	SovCISS	Encompassing
2008 crisis dummy	0.041**		0.039**
	(0.016)		(0.014)
Debt crisis dummy	0.010		0.010
	(0.009)		(0.009)
COVID dummy	−0.004		−0.002
	(0.012)		(0.012)
GDP growth	−0.001	−0.002*	−0.001
	(0.001)	(0.001)	(0.001)
Public debt (% GDP)	−0.000	−0.001	−0.000
	(0.001)	(0.001)	(0.001)
Unemployment rate	−0.000	−0.000	−0.001
	(0.001)	(0.001)	(0.001)
SovCISS (mean)		0.039	0.018
		(0.035)	(0.034)
Num.Obs.	209	209	209
R2	0.919	0.913	0.920
R2 Within	0.159	0.096	0.162
AIC	−684.7	−673.7	−683.6
BIC	−627.9	−623.5	−623.4

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs. Country fixed effects only (crisis dummies collinear with year FE).

Table 9 - Estimator comparison: pooled slope, heterogeneous slopes, common factors

Estimator	Coefficient	SE	p-value
Two−way FE (baseline)	−0.0249	0.0615	0.695
Mean Group	+0.0194	0.0210	0.355
CCEMG (common factors)	−0.0215	0.0377	0.568
MG slope dispersion (SovCISS)	range [-0.0783; +0.1561]	sd=0.0697	6/11 neg.
Chow poolability (F)	F=5.90		0.000
First differences (β, d_sov)	-0.0286	0.0316	0.387
Anderson-Hsiao IV (β)	-0.0320	0.0562	first-stage F=39.3
GMM instrument count	153 (N=11)		

Coefficient on SovCISS in the H1 specification. The two-way FE estimator imposes a common slope; the Mean Group estimator averages country-specific slopes; the CCEMG estimator additionally proxies unobserved common factors through cross-sectional averages, allowing heterogeneous exposure to shocks common to the panel. With T = 19, MG estimates 5 parameters per country and CCEMG 10; p-values are two-sided.

Table 10a - Extended sample: eurozone and non-eurozone EU countries

Group	N	Cult. exp (% GDP)	SovCISS (mean)	SovCISS (max)	Debt (% GDP)	GDP growth (%)	Unemployment (%)
Eurozone (11)	209	0.425	0.218	0.946	92.255	1.329	9.038
Non-eurozone (5)	95	0.684	0.174	0.627	48.317	2.155	6.276

Notes:

Sample means, 2006-2024. Non-eurozone group: Czechia, Denmark, Hungary, Poland, Sweden. Sources as in Table 1; SovCISS from the ECB Data Portal, which publishes the indicator for euro area and non-euro area EU countries.

Table 10b - Slope homogeneity across monetary regimes

	Pooled (16)	+ Interaction	Fully interacted	Eurozone (11)	Non-eurozone (5)
SovCISS (mean)	0.019	−0.176	−0.168	−0.025	−0.130
	(0.055)	(0.104)	(0.113)	(0.062)	(0.171)
GDP growth	0.001	0.001	0.006	−0.002	0.023
	(0.004)	(0.003)	(0.006)	(0.001)	(0.014)
Public debt (% GDP)	0.000	0.000	0.003	0.000	0.003
	(0.001)	(0.001)	(0.004)	(0.001)	(0.003)
Unemployment rate	−0.005	−0.006	−0.021	−0.000	−0.020
	(0.006)	(0.006)	(0.019)	(0.004)	(0.025)
SovCISS x Eurozone		0.232	0.149		
		(0.149)	(0.102)		
GDP growth x Eurozone			−0.004		
			(0.004)		
Public debt x Eurozone			−0.003		
			(0.004)		
Unemployment x Eurozone			0.021		
			(0.020)		
Num.Obs.	304	304	304	209	95
R2	0.871	0.878	0.890	0.929	0.656
R2 Within	0.024	0.078	0.165	0.014	0.192
Chow F (partition, 4 df)	0.99				
p (partition)	0.444				
MG slope range (16 countries)	[-0.240; +0.156]				

* p < 0.1, ** p < 0.05, *** p < 0.01

Country-clustered SEs. Country and year fixed effects in all models. The eurozone indicator is time-invariant and absorbed by the country fixed effects, so only its interactions are identified. The non-eurozone column has five clusters and its SEs are indicative only. Stars denote two-sided significance.

Table 10c — Sequential poolability of the SovCISS slope

Step	Excluded	F stat	p value
1	Ireland	4.013	< 0.0001
2	Portugal	3.100	0.0019
stop	— (null not rejected)	0.875	0.5394
1	Ireland	4.606	< 0.0001
2	Portugal	3.436	< 0.0001
3	Hungary	2.428	0.0044
stop	— (null not rejected)	1.137	0.3320

Panel A — 11 euro-area countries (final group: 9)

Panel B — 16 EU countries (final group: 13)

Note.

Sequential exclusion following Kapetanios. At each step the joint Wald test of equality of the country-specific SovCISS slopes is computed; the country whose removal lowers the statistic most is excluded. Heteroskedasticity-robust variance (clustering by country is infeasible with one interaction per cluster). Controls: GDP growth, public debt, unemployment; country and year fixed effects. Panel B poolable group (13 countries). Euro area: Austria, Belgium, Finland, France, Germany, Greece, Italy, Netherlands, Spain; outside: Czechia, Denmark, Poland, Sweden. The procedure does not correct for multiple testing.

Table 10d - Every specification re-estimated on the extended sample

Specification	Eurozone (11)	Extended (16)	N_11	N_16
H1 Cultural exp. (% GDP), SovCISS	−0.0249 (0.0615)	+0.0187 (0.0550)	209	304
H2 Cultural exp. (% total exp.), SovCISS	−0.0365 (0.1096)	+0.0381 (0.1185)	209	304
H4 Lagged SovCISS (t−1)	−0.0261 (0.0457)	+0.0916 (0.0813)	198	288
H5 Quadratic term SovCISS squared	−0.0795 (0.1169)	+0.0319 (0.1383)	209	304
R5 log real expenditure per capita	−0.1875 (0.1823)	−0.0343 (0.1448)	209	304
CCEMG (common factors), SovCISS	−0.0215 (0.0377)	−0.0761* (0.0460)	209	304

Notes: Coefficient of interest with country-clustered SE in parentheses. Country and year fixed effects throughout. The extended sample adds Czechia, Denmark, Hungary, Poland and Sweden, for which the ECB publishes the SovCISS. H3 is not re-estimated: the WGI-GE moderator is unavailable for the five added countries. Stars denote two-sided significance. * p<0.1, ** p<0.05, *** p<0.01.

Appendix - R code

```
## ======================================================================
## APPENDIX — R CODE
## Estimation, testing and simulation code. Table and figure rendering
## code is omitted; the tables themselves are reproduced in the annex.
## ======================================================================

## ======================================================================
## 1. Parlgov.R — Construction of the annual government-ideology series
(ParlGov)
##    Input series for robustness check R4
## ======================================================================

library(readxl)
library(dplyr)
library(tidyr)
library(writexl)

## Paths
in_path  <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/parlgov-
stable.xlsx"
out_path <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/
ParlGov_annual_ideology.xlsx"

target_iso3 <- c("AUT","BEL","DEU","ESP","FIN","FRA","GRC",
                 "IRL","ITA","NLD","PRT")

## 1. Read
cab <- read_excel(in_path, sheet = "cabinet", guess_max = 200000)
stopifnot(all(c("country_name_short","start_date","cabinet_name",
                "cabinet_party","seats","left_right") %in% names(cab)))

## 2. Clean and filter
if (!is.numeric(cab$left_right))
  cab$left_right <- as.numeric(gsub(",", ".", cab$left_right))

to_date_safe <- function(x) {
  if (inherits(x, "Date"))   return(x)
  if (inherits(x, "POSIXt")) return(as.Date(x))
  if (is.numeric(x))         return(as.Date(x, origin = "1899-12-30"))
  out <- rep(as.Date(NA), length(x))
  x   <- as.character(x)
  ok1 <- !is.na(x) & grepl("^\\d{4}-\\d{2}-\\d{2}", x)
  out[ok1] <- as.Date(substr(x[ok1], 1, 10), format = "%Y-%m-%d")
  ok2 <- !is.na(x) & is.na(out) & grepl("^\\d{2}/\\d{2}/\\d{4}$", x)
  out[ok2] <- as.Date(x[ok2], format = "%d/%m/%Y")
  ok3 <- !is.na(x) & is.na(out) & grepl("^\\d+(\\.\\d+)?$", x)
  out[ok3] <- as.Date(as.numeric(x[ok3]), origin = "1899-12-30")
  out
}

cab$start_date <- to_date_safe(cab$start_date)
n_bad <- sum(is.na(cab$start_date) & !is.na(cab$country_name_short))
```

```r
df <- cab |>
  filter(country_name_short %in% target_iso3,
         cabinet_party == 1L,
         !is.na(left_right),
         !is.na(seats),
         !is.na(start_date),
         seats > 0) |>
  select(country = country_name_short, start_date, cabinet_name,
         seats, left_right)

## 3. Cabinet-level seat-weighted ideology score
cab_ideo <- df |>
  group_by(country, start_date, cabinet_name) |>
  summarise(cab_lr  = weighted.mean(left_right, seats),
            n_parties = n(),
            total_seats = sum(seats),
            .groups = "drop") |>
  arrange(country, start_date)

## 4. Derive each cabinet's end date
cab_ideo <- cab_ideo |>
  group_by(country) |>
  arrange(start_date, .by_group = TRUE) |>
  mutate(end_date = dplyr::lead(start_date) - 1L,
         end_date = if_else(is.na(end_date), as.Date("2099-12-31"),
end_date)) |>
  ungroup()

## 5. Coverage report: last cabinet per country
last_cab <- cab_ideo |>
  group_by(country) |>
  summarise(last_cabinet_start = max(start_date),
            last_cabinet_name  = cabinet_name[which.max(start_date)],
            .groups = "drop") |>
  arrange(last_cabinet_start)

print(as.data.frame(last_cab))

## 6. Annual mean weighted by days in office
years <- 2006:2024

expand_year <- function(y, cab_country) {
  y_start <- as.Date(paste0(y, "-01-01"))
  y_end   <- as.Date(paste0(y, "-12-31"))
  cabs <- cab_country |>
    filter(start_date <= y_end, end_date >= y_start) |>
    mutate(overlap_start = pmax(start_date, y_start),
           overlap_end   = pmin(end_date,   y_end),
           days = as.integer(overlap_end - overlap_start) + 1L) |>
    filter(days > 0)
  if (nrow(cabs) == 0) return(NA_real_)
  weighted.mean(cabs$cab_lr, cabs$days)
}

annual <- expand_grid(country = target_iso3, year = years) |>
  group_by(country) |>
```

```r
  mutate(gov_lr = vapply(year,
                         function(y) expand_year(y,
cab_ideo[cab_ideo$country == country[1], ]),
                         numeric(1))) |>
  ungroup()

## 7. Flag country-years extrapolated past the last cabinet
annual <- annual |>
  left_join(last_cab |> select(country, last_cabinet_start),
            by = "country") |>
  mutate(flag_extrapolated =
           as.Date(paste0(year, "-12-31")) > last_cabinet_start + 365L) |
>
  select(-last_cabinet_start)

extrap <- annual |> filter(flag_extrapolated)
if (nrow(extrap) > 0) print(as.data.frame(extrap))

## 8. Wide layout
wide <- annual |>
  select(country, year, gov_lr) |>
  pivot_wider(names_from = year, values_from = gov_lr) |>
  arrange(country)

print(as.data.frame(wide), digits = 4)

## 9. Write to xlsx
write_xlsx(list(wide      = wide,
                long      = annual,
                cabinets = cab_ideo,
                coverage = last_cab),
           out_path)


## ==================================================================
## 2. H1H2V2.R — H1 and H2: baseline specification
##    Descriptive statistics, collinearity diagnostics, nested models,
bootstrap, MDE
## ==================================================================

## Packages
library(readxl); library(dplyr); library(tidyr)
library(fixest); library(plm)

## §1 — Import via header patterns (standard block, shared with H3/H4/H5)

path <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données
groupées.xlsx"
df_raw <- read_excel(path, sheet = "Données")
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]

h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))

pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf(
```

```r
    "Motif '%s' -> %d colonne(s) [%s]. Corrige ce motif.",
    pattern, length(idx), paste(names(df_raw)[idx], collapse = " | ")))
  idx
}

map <- c(
  country                  = "^country$",
  year                     = "^time$|^year$",
  cult_exp_meur            = "expenditure \\(m€\\)$",
  cult_exp_gdp             = "cultural services exp \\(%gdp\\)",
  tot_exp_meur             = "^total exp \\(m€\\)",
  cult_exp_totexp          = "cult serv exp \\(%tot",
  debt_gdp                 = "debt stock",
  inflation                = "^inflation",
  gdp_growth               = "growth rate",
  unemployment             = "unemployment",
  sovciss_mean             = "sovciss.*annual",
  sovciss_max              = "sovciss.*max",
  eurozone                 = "^eurozone$"
)

sel <- sapply(map, pick)
df  <- df_raw[, sel]; names(df) <- names(map)

df <- df %>%
  mutate(country = trimws(country)) %>%
  fill(country, .direction = "down") %>%
  filter(!is.na(year)) %>%
  mutate(country = as.factor(country),
         year     = as.integer(year),
         eurozone = as.integer(eurozone)) %>%
  filter(eurozone == 1) %>%
  mutate(country = droplevels(country))

stopifnot(nrow(df) == 209, length(unique(df$country)) == 11)
stopifnot(all(df$sovciss_mean >= 0 & df$sovciss_mean <= 1, na.rm = TRUE))

ctrl <- "gdp_growth + debt_gdp + unemployment"

## §1bis — Descriptive statistics (Table 1)

vars_desc <- c(
  "cult_exp_gdp"     = "Cultural exp. (% GDP)",
  "cult_exp_totexp" = "Cultural exp. (% total gov. exp.)",
  "sovciss_mean"     = "SovCISS (annual mean)",
  "gdp_growth"       = "Real GDP growth (%)",
  "debt_gdp"         = "Public debt (% GDP)",
  "unemployment"     = "Unemployment rate (%)"
)

desc_stats <- df %>%
  select(all_of(names(vars_desc))) %>%
  summarise(across(everything(),
                   list(N       = ~sum(!is.na(.)),
                        Mean    = ~mean(., na.rm = TRUE),
                        SD      = ~sd(., na.rm = TRUE),
```

```r
                              Min      = ~min(., na.rm = TRUE),
                              P25      = ~quantile(., 0.25, na.rm = TRUE),
                              Median   = ~median(., na.rm = TRUE),
                              P75      = ~quantile(., 0.75, na.rm = TRUE),
                              Max      = ~max(., na.rm = TRUE)),
                        .names = "{.col}__{.fn}")) %>%
  tidyr::pivot_longer(everything(),
                            names_to = c("Variable", "Statistic"),
                            names_sep = "__") %>%
  tidyr::pivot_wider(names_from = Statistic, values_from = value) %>%
  mutate(Variable = vars_desc[Variable],
         across(where(is.numeric), ~round(., 3)))

print(as.data.frame(desc_stats))

## §2 — Multicollinearity: correlations and VIF (Tables 2a, 2b)

cor_mat <- cor(df[, c("sovciss_mean", "gdp_growth", "debt_gdp",
"unemployment")],
               use = "complete.obs")
print(round(cor_mat, 3))

cor_df <- as.data.frame(round(cor_mat, 3))
cor_df <- cbind(Variable = rownames(cor_df), cor_df)
rownames(cor_df) <- NULL
labels_map <- c("sovciss_mean"  = "SovCISS (mean)",
                "gdp_growth"    = "GDP growth",
                "debt_gdp"      = "Public debt (% GDP)",
                "unemployment"  = "Unemployment rate")
cor_df$Variable <- labels_map[cor_df$Variable]
names(cor_df)[-1] <- labels_map[names(cor_df)[-1]]

vif_ols <- lm(cult_exp_gdp ~ sovciss_mean + gdp_growth + debt_gdp +
unemployment,
              data = df)

if (requireNamespace("car", quietly = TRUE)) {
  vif_vec <- car::vif(vif_ols)
  print(round(vif_vec, 2))

  vif_df <- data.frame(Variable = labels_map[names(vif_vec)],
                       VIF      = round(as.numeric(vif_vec), 2))

}
## §3 — Inference diagnostics: Pesaran CD and Wooldridge

pdf <- pdata.frame(df, index = c("country", "year"))
m_plm <- plm(cult_exp_gdp ~ sovciss_mean + gdp_growth + debt_gdp +
unemployment,
             data = pdf, model = "within", effect = "twoways")

print(pcdtest(m_plm, test = "cd"))
print(pwartest(m_plm))

## §3ter — Pesaran CD with country FE only vs two-way FE; average
pairwise residual correlation
```

```r
library(plm)

pdat <- pdata.frame(df, index = c("country", "year"))
f_h1 <- cult_exp_gdp ~ sovciss_mean + gdp_growth + debt_gdp +
unemployment

m_one <- plm(f_h1, data = pdat, model = "within", effect = "individual")
m_two <- plm(f_h1, data = pdat, model = "within", effect = "twoways")

cd_one <- pcdtest(m_one, test = "cd")
cd_two <- pcdtest(m_two, test = "cd")

avg_pairwise_corr <- function(mod) {
  idx <- index(mod)
  w <- data.frame(country = as.character(idx[[1]]),
                  year    = as.character(idx[[2]]),
                  e       = as.numeric(residuals(mod)),
                  stringsAsFactors = FALSE)
  M <- tapply(w$e, list(w$year, w$country), identity)
  C <- cor(M, use = "pairwise.complete.obs")
  mean(C[upper.tri(C)], na.rm = TRUE)
}

N <- length(unique(as.character(index(m_two)[[1]])))

idx <- index(m_two)
d <- data.frame(country = as.character(idx[[1]]),
                year    = as.numeric(as.character(idx[[2]])),
                e       = as.numeric(residuals(m_two)),
                stringsAsFactors = FALSE)
d <- d[order(d$country, d$year), ]
d$e_lag <- ave(d$e, d$country, FUN = function(z) c(NA, head(z, -1)))
ar1 <- lm(e ~ e_lag - 1, data = d)

print(pwartest(f_h1, data = pdat))

## §3bis — Identifying variation of SovCISS: variance decomposition

v_country <- summary(lm(sovciss_mean ~ factor(country),
data = df))$r.squared
v_year    <- summary(lm(sovciss_mean ~ factor(year),
data = df))$r.squared
fe_dec    <- lm(sovciss_mean ~ factor(country) + factor(year),
data = df)
v_both    <- summary(fe_dec)$r.squared
sd_raw    <- sd(df$sovciss_mean, na.rm = TRUE)
sd_res    <- sd(resid(fe_dec))

## §4 — Nested models M1-M3, H1 then H2 (Table 3a)

h1_m1 <- feols(cult_exp_gdp ~ sovciss_mean | country + year,
               data = df, cluster = ~country)
h1_m2 <- feols(cult_exp_gdp ~ sovciss_mean + gdp_growth | country + year,
               data = df, cluster = ~country)
h1_m3 <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl, "|
country + year")),
```

```
                  data = df, cluster = ~country)

h2_m1 <- feols(cult_exp_totexp ~ sovciss_mean | country + year,
                  data = df, cluster = ~country)
h2_m2 <- feols(cult_exp_totexp ~ sovciss_mean + gdp_growth | country +
year,
                  data = df, cluster = ~country)
h2_m3 <- feols(as.formula(paste("cult_exp_totexp ~ sovciss_mean +", ctrl,
"| country + year")),
                  data = df, cluster = ~country)

print(etable(h1_m1, h1_m2, h1_m3,
                headers = c("M1: SovCISS only", "M2: + growth", "M3:
baseline"),
                tex = FALSE))

print(etable(h2_m1, h2_m2, h2_m3,
                headers = c("M1: SovCISS only", "M2: + growth", "M3:
baseline"),
                tex = FALSE))

h1_dk <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl, "|
country + year")),
                  data = df, panel.id = ~country + year, vcov = DK(2))
h2_dk <- feols(as.formula(paste("cult_exp_totexp ~ sovciss_mean +", ctrl,
"| country + year")),
                  data = df, panel.id = ~country + year, vcov = DK(2))

## §5 — Complementary diagnostics (no-year-FE bound; log real
expenditure)

h1_no_yfe <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +",
ctrl, "| country")),
                       data = df, cluster = ~country)

print(etable(h1_m3, h1_no_yfe, headers = c("Baseline", "Sans FE annee"),
tex = FALSE))

if ("cult_exp_meur" %in% names(df) && "inflation" %in% names(df)) {
  df <- df %>% group_by(country) %>%
    mutate(cpi = cumprod(1 + coalesce(inflation, 0) / 100),
           cult_exp_real = cult_exp_meur / cpi) %>% ungroup() %>%
    mutate(log_cult_real = log(cult_exp_real))

  h1_log <- feols(as.formula(paste("log_cult_real ~ sovciss_mean +",
ctrl, "| country + year")),
                     data = df, cluster = ~country)

  print(etable(h1_log, headers = "log(cultural exp, real)", tex = FALSE))
}

## §6 — Wild cluster bootstrap by exact enumeration (2^11 = 2,048
vectors)

boot_enum <- function(data, dep, rhs_full, rhs_restr, param) {
  fe <- "| country + year"
```

```
  m_r <- feols(as.formula(paste(dep, "~", rhs_restr, fe)), data = data)
  data$fitted_r <- fitted(m_r); data$resid_r <- resid(m_r)

  m_f <- feols(as.formula(paste(dep, "~", rhs_full, fe)), data = data,
cluster = ~country)
  t_obs <- coeftable(m_f)[param, "t value"]

  cl <- sort(unique(as.character(data$country)))
  W  <- as.matrix(expand.grid(rep(list(c(-1, 1)), length(cl))))
  fml_b <- as.formula(paste("y_star ~", rhs_full, fe))

  t_boot <- apply(W, 1, function(w) {
    names(w) <- cl
    data$y_star <- data$fitted_r + data$resid_r *
w[as.character(data$country)]
    coeftable(feols(fml_b, data = data, cluster = ~country))[param, "t
value"]
  })
  list(t_obs = t_obs,
       p_two  = mean(abs(t_boot) >= abs(t_obs)),
       p_less = mean(t_boot <= t_obs))
}

rhs_full <- paste("sovciss_mean +", ctrl)
boot_h1 <- boot_enum(df, "cult_exp_gdp",    rhs_full, ctrl,
"sovciss_mean")
boot_h2 <- boot_enum(df, "cult_exp_totexp", rhs_full, ctrl,
"sovciss_mean")

## §7 — Driscoll-Kraay sensitivity to the truncation lag

for (L in 1:3) {
  m_dk <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
"| country + year")),
                data = df, panel.id = ~country + year, vcov = DK(L))

}

## §8 — Minimum detectable effect (MDE)

mde_from <- function(m, param, power = 0.80, df_t = 10) {
  se <- sqrt(vcov(m)[param, param])
  (qt(0.975, df_t) + qt(power, df_t)) * se
}

amp <- df %>% group_by(country) %>%
  summarise(a = max(sovciss_mean, na.rm = TRUE) - min(sovciss_mean, na.rm
= TRUE))
delta_sov <- median(amp$a)
mean_dv_h1 <- mean(df$cult_exp_gdp, na.rm = TRUE)
mean_dv_h2 <- mean(df$cult_exp_totexp, na.rm = TRUE)

mde_h1 <- mde_from(h1_m3, "sovciss_mean")
mde_h2 <- mde_from(h2_m3, "sovciss_mean")

b1  <- coef(h1_m3)["sovciss_mean"]
```

```r
se1 <- sqrt(vcov(h1_m3)["sovciss_mean", "sovciss_mean"])
ci1 <- b1 + c(-1, 1) * qt(0.975, 10) * se1


## =====================================================================
## 3. H3V3.R — H3: institutional quality as moderator of financial stress
##    Interaction models, marginal effects, EU-IMF programme control,
bootstrap, MDE
## =====================================================================

## Packages
library(readxl); library(dplyr); library(tidyr)
library(fixest); library(car)

## §1 — Import via header patterns (standard block)

path <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données
groupées.xlsx"
df_raw <- read_excel(path, sheet = "Données")
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]

h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))

pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf(
    "Motif '%s' -> %d colonne(s) [%s]. Corrige ce motif.",
    pattern, length(idx), paste(names(df_raw)[idx], collapse = " | ")))
  idx
}

map <- c(
  country           = "^country$",
  year              = "^time$|^year$",
  cult_exp_gdp      = "cultural services exp \\(%gdp\\)",
  cult_exp_totexp   = "cult serv exp \\(%tot",
  debt_gdp          = "debt stock",
  gdp_growth        = "growth rate",
  unemployment      = "unemployment",
  sovciss_mean      = "sovciss.*annual",
  wgi_ge            = "wgi",
  eurozone          = "^eurozone$"
)

sel <- sapply(map, pick)
df  <- df_raw[, sel]; names(df) <- names(map)

df <- df %>%
  mutate(country = trimws(country)) %>%
  fill(country, .direction = "down") %>%
  filter(!is.na(year)) %>%
  mutate(country = as.factor(country),
         year      = as.integer(year),
         eurozone = as.integer(eurozone)) %>%
  filter(eurozone == 1) %>%
  mutate(country = droplevels(country))
```

```
stopifnot(nrow(df) == 209, length(unique(df$country)) == 11)

## §2 — Derived variables: centred terms and EU-IMF programme dummy

df <- df %>% mutate(
  sovciss_c = sovciss_mean - mean(sovciss_mean, na.rm = TRUE),
  wgi_c     = wgi_ge       - mean(wgi_ge,       na.rm = TRUE),
  troika    = as.integer(country %in% c("Greece", "Ireland", "Portugal"))
)

df$south_troika <- NULL

r_gw <- cor(df$troika, df$wgi_c, use = "complete.obs")

ctrl <- "gdp_growth + debt_gdp + unemployment"

## §3 — Table 4a: interaction only, + WGI-GE main effect, + controls

h3_bare <- feols(cult_exp_gdp ~ sovciss_c + sovciss_c:wgi_c
                 | country + year, data = df, cluster = ~country)

h3_wgi  <- feols(cult_exp_gdp ~ sovciss_c * wgi_c
                 | country + year, data = df, cluster = ~country)

h3_full <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c * wgi_c +",
ctrl,
                                  "| country + year")),
                 data = df, cluster = ~country)

h3_full_dk <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c * wgi_c
+", ctrl,
                                     "| country + year")),
                    data = df, panel.id = ~country + year, vcov = DK(2))

print(etable(h3_bare, h3_wgi, h3_full,
             headers = c("Interaction only", "+ WGI main", "M3
baseline")))

## §4 — Table 4b: marginal effect of stress by institutional quality

qs   <- c(.10, .50, .90)
pcts <- quantile(df$wgi_ge, qs)
lab  <- c("10th (weakest)", "50th (median)", "90th (strongest)")

est <- se <- tv <- pv <- numeric(length(qs))
for (i in seq_along(qs)) {
  wc     <- pcts[i] - mean(df$wgi_ge)
  me     <- deltaMethod(h3_full,
                        sprintf("sovciss_c + %.6f*`sovciss_c:wgi_c`",
wc),
                        vcov = vcov(h3_full))
  est[i] <- me$Estimate
  se[i]  <- me$SE
  tv[i]  <- est[i] / se[i]
  pv[i]  <- 2 * pt(-abs(tv[i]), df = 10)
```

```r
}
star <- ifelse(pv < .01, "***", ifelse(pv < .05, "**", ifelse(pv < .1,
"*", "")))

me_tab <- data.frame(
  "Percentile of WGI-GE" = lab,
  "WGI-GE"               = sprintf("%.2f", pcts),
  "Marginal effect"      = sprintf("%+.3f%s", est, star),
  "SE"                   = sprintf("%.3f", se),
  "t"                    = sprintf("%+.2f", tv),
  "p-value"              = sprintf("%.3f", pv),
  check.names = FALSE, stringsAsFactors = FALSE)

me_tab[] <- lapply(me_tab, function(x) gsub("-", "\u2212", x, fixed =
TRUE))

print(me_tab)

## §5 — Table 4c: control for the EU-IMF programme interaction

h3_troika <- feols(as.formula(paste(
  "cult_exp_gdp ~ sovciss_c * wgi_c + sovciss_c:troika +", ctrl,
  "| country + year")),
  data = df, cluster = ~country)

print(etable(h3_full, h3_troika,
             headers = c("Baseline", "+ SovCISS x EU-IMF programme")))

## §5bis — Table 4d: implied country-level marginal effects

cty <- df %>%
  group_by(country) %>%
  summarise(wgi_mean   = mean(wgi_ge, na.rm = TRUE),
            wgi_c_mean = mean(wgi_c,  na.rm = TRUE),
            prog       = max(troika),
            .groups = "drop") %>%
  arrange(wgi_mean)

est <- se <- tv <- pv <- numeric(nrow(cty))
for (i in seq_len(nrow(cty))) {
  expr <- sprintf("sovciss_c + %.6f*`sovciss_c:wgi_c`",
cty$wgi_c_mean[i])
  if (cty$prog[i] == 1) expr <- paste0(expr, " + `sovciss_c:troika`")
  me     <- deltaMethod(h3_troika, expr, vcov = vcov(h3_troika))
  est[i] <- me$Estimate
  se[i]  <- me$SE
  tv[i]  <- est[i] / se[i]
  pv[i]  <- 2 * pt(-abs(tv[i]), df = 10)
}
star <- ifelse(pv < .01, "***", ifelse(pv < .05, "**", ifelse(pv < .1,
"*", "")))

me4d <- data.frame(
  "Country"                  = as.character(cty$country),
  "WGI-GE (mean)"            = sprintf("%.2f", cty$wgi_mean),
  "EU-IMF programme"         = ifelse(cty$prog == 1, "Yes", "No"),
```

```r
  "Implied marginal effect" = sprintf("%+.3f%s", est, star),
  "SE"                      = sprintf("%.3f", se),
  "p-value"                 = sprintf("%.3f", pv),
  check.names = FALSE, stringsAsFactors = FALSE)

me4d[] <- lapply(me4d, function(x) gsub("-", "\u2212", x, fixed = TRUE))

print(me4d)

## §6 — Enumerated wild cluster bootstrap on gamma_3

m_r <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c + wgi_c +", ctrl,
                              "| country + year")),
             data = df)
df$fitted_r <- fitted(m_r); df$resid_r <- resid(m_r)

t_obs <- coeftable(h3_full)["sovciss_c:wgi_c", "t value"]

cl <- sort(unique(as.character(df$country)))
W  <- as.matrix(expand.grid(rep(list(c(-1, 1)), length(cl))))

fml_boot <- as.formula(paste("y_star ~ sovciss_c * wgi_c +", ctrl,
                             "| country + year"))

t_boot <- apply(W, 1, function(w) {
  names(w) <- cl
  df$y_star <- df$fitted_r + df$resid_r * w[as.character(df$country)]
  coeftable(feols(fml_boot, data = df, cluster = ~country))
["sovciss_c:wgi_c", "t value"]
})

p_two     <- mean(abs(t_boot) >= abs(t_obs))
p_greater <- mean(t_boot >= t_obs)

## §6ter — Enumerated bootstrap on both interactions of Table 4c

boot_enum_4c <- function(rhs_full, rhs_restr, param, data = df,
                         dep = "cult_exp_gdp") {
  m_r <- feols(as.formula(paste(dep, "~", rhs_restr, "| country +
year")),
               data = data)
  data$fitted_r <- fitted(m_r)
  data$resid_r  <- resid(m_r)

  m_u   <- feols(as.formula(paste(dep, "~", rhs_full, "| country +
year")),
                 data = data, cluster = ~country)
  t_obs <- coeftable(m_u)[param, "t value"]

  cl <- sort(unique(as.character(data$country)))
  W  <- as.matrix(expand.grid(rep(list(c(-1, 1)), length(cl))))

  fml_b <- as.formula(paste("y_star ~", rhs_full, "| country + year"))

  t_boot <- apply(W, 1, function(w) {
    names(w) <- cl
```

```
    data$y_star <- data$fitted_r + data$resid_r *
w[as.character(data$country)]
    coeftable(feols(fml_b, data = data, cluster = ~country))[param, "t
value"]
  })

  list(t_obs = t_obs, p_two = mean(abs(t_boot) >= abs(t_obs)))
}

rhs_4c <- paste("sovciss_c * wgi_c + sovciss_c:troika +", ctrl)

b_int  <- boot_enum_4c(
  rhs_4c,
  paste("sovciss_c + wgi_c + sovciss_c:troika +", ctrl),
  "sovciss_c:wgi_c")

b_prog <- boot_enum_4c(
  rhs_4c,
  paste("sovciss_c * wgi_c +", ctrl),
  "sovciss_c:troika")

## §7 — Driscoll-Kraay sensitivity on the interaction

for (L in 1:3) {
  m_dk <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c * wgi_c +",
ctrl,
                                "| country + year")),
                data = df, panel.id = ~country + year, vcov = DK(L))

}

## §8 — MDE on gamma_3

mde_from <- function(m, param, power = 0.80, df_t = 10) {
  se <- sqrt(vcov(m)[param, param])
  (qt(0.975, df_t) + qt(power, df_t)) * se
}

mde_g3 <- mde_from(h3_full, "sovciss_c:wgi_c")
gamma3 <- coef(h3_full)["sovciss_c:wgi_c"]


## ======================================================================
## 4. H4V2.R — H4: persistence of the stress effect
##    Lag models, cumulative effects, bootstrap, MDE
## ======================================================================

## Packages
library(readxl); library(dplyr); library(tidyr)
library(fixest); library(car)

## §1 — Import via header patterns

path <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données
groupées.xlsx"
df_raw <- read_excel(path, sheet = "Données")
```

```r
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]

h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))

pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf(
    "Motif '%s' -> %d colonne(s) [%s]. Corrige ce motif.",
    pattern, length(idx), paste(names(df_raw)[idx], collapse = " | ")))
  idx
}

map <- c(
  country      = "^country$",
  year         = "^time$|^year$",
  cult_exp_gdp = "cultural services exp \\(%gdp\\)",
  debt_gdp     = "debt stock",
  gdp_growth   = "growth rate",
  unemployment = "unemployment",
  sovciss_mean = "sovciss.*annual",
  eurozone     = "^eurozone$"
)

sel <- sapply(map, pick)
df  <- df_raw[, sel]; names(df) <- names(map)

df <- df %>%
  mutate(country = trimws(country)) %>%
  fill(country, .direction = "down") %>%
  filter(!is.na(year)) %>%
  mutate(country = as.factor(country),
         year     = as.integer(year),
         eurozone = as.integer(eurozone)) %>%
  filter(eurozone == 1) %>%
  mutate(country = droplevels(country))

stopifnot(nrow(df) == 209, length(unique(df$country)) == 11)

## §2 — Derived variables: lags 1 and 2

df <- df %>%
  arrange(country, year) %>%
  group_by(country) %>%
  mutate(
    sovciss_lag1 = dplyr::lag(sovciss_mean, 1),
    sovciss_lag2 = dplyr::lag(sovciss_mean, 2)
  ) %>%
  ungroup()

stopifnot(sum(is.na(df$sovciss_lag1)) == 11)
stopifnot(sum(is.na(df$sovciss_lag2)) == 22)

df_lag <- filter(df, !is.na(sovciss_lag1)); stopifnot(nrow(df_lag) ==
198)
df_dl  <- filter(df, !is.na(sovciss_lag2)); stopifnot(nrow(df_dl)  ==
187)
```

```r
ctrl <- "gdp_growth + debt_gdp + unemployment"

## §3 — Nested models: contemporaneous, lag only, joint

h4_contemp <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
                                     "| country + year")),
                    data = df_lag, cluster = ~country)

h4_lag_only <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_lag1 +", ctrl,
                                      "| country + year")),
                     data = df_lag, cluster = ~country)

h4_joint <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean + sovciss_lag1 +",
                                   ctrl, "| country + year")),
                  data = df_lag, cluster = ~country)

h4_joint_dk <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean + sovciss_lag1 +",
                                      ctrl, "| country + year")),
                     data = df_lag, panel.id = ~country + year, vcov = DK(2))

print(etable(h4_contemp, h4_lag_only, h4_joint, h4_joint_dk,
             headers = c("Contemporain", "Lagge seul", "Joint", "Joint DK(2)"),
             tex = FALSE))

corr_lag <- cor(df_lag$sovciss_mean, df_lag$sovciss_lag1, use = "complete.obs")

## §4 — Two-year cumulative effect (delta method) and joint Wald test

print(deltaMethod(h4_joint, "sovciss_mean + sovciss_lag1", vcov = vcov(h4_joint)))

print(wald(h4_joint, keep = "sovciss"))

## §5 — Extension: second lag and three-year cumulative effect

h4_dl <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean + sovciss_lag1 + sovciss_lag2 +",
                                ctrl, "| country + year")),
               data = df_dl, cluster = ~country)

print(etable(h4_dl, headers = "SovCISS(t) + t-1 + t-2", tex = FALSE))

print(deltaMethod(h4_dl,
                  "sovciss_mean + sovciss_lag1 + sovciss_lag2",
                  vcov = vcov(h4_dl)))

print(wald(h4_dl, keep = "sovciss"))
```

```r
## §4bis — Cumulative-effect rows added to Table 5

dm2 <- deltaMethod(h4_joint, "sovciss_mean + sovciss_lag1", vcov =
vcov(h4_joint))
w2  <- wald(h4_joint, keep = "sovciss")

dm3 <- deltaMethod(h4_dl, "sovciss_mean + sovciss_lag1 + sovciss_lag2",
vcov = vcov(h4_dl))
w3  <- wald(h4_dl, keep = "sovciss")

fmt_ci <- function(dm) sprintf("[%.4f ; %.4f]", dm$`2.5 %`, dm$`97.5 %`)

cum_rows <- data.frame(
  term = c("Cumulative effect (2-year)", "95% CI",
           "Cumulative effect (3-year)", "95% CI",
           "Joint Wald (2-year, F)", "Joint Wald (3-year, F)"),
  Contemporaneous = "",
  `Lag only` = "",
  Joint = c(sprintf("%+.4f", dm2$Estimate), fmt_ci(dm2), "", "",
            sprintf("F=%.2f, p=%.3f", w2[["stat"]], w2[["p"]]), ""),
  `Joint DK(2)` = "",
  `Two lags (n=187)` = c("", "", sprintf("%+.4f", dm3$Estimate),
fmt_ci(dm3),
                         "", sprintf("F=%.2f, p=%.3f", w3[["stat"]],
w3[["p"]])),
  check.names = FALSE, stringsAsFactors = FALSE
)

## §6 — Enumerated wild cluster bootstrap on gamma_4

m_r <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
                              "| country + year")),
             data = df_lag)
df_lag$fitted_r <- fitted(m_r); df_lag$resid_r <- resid(m_r)

t_obs <- coeftable(h4_joint)["sovciss_lag1", "t value"]

cl <- sort(unique(as.character(df_lag$country)))
W  <- as.matrix(expand.grid(rep(list(c(-1, 1)), length(cl))))

fml_boot <- as.formula(paste("y_star ~ sovciss_mean + sovciss_lag1 +",
                             ctrl, "| country + year"))

t_boot <- apply(W, 1, function(w) {
  names(w) <- cl
  df_lag$y_star <- df_lag$fitted_r + df_lag$resid_r *
w[as.character(df_lag$country)]
  coeftable(feols(fml_boot, data = df_lag, cluster = ~country))
["sovciss_lag1", "t value"]
})

p_two  <- mean(abs(t_boot) >= abs(t_obs))
p_less <- mean(t_boot <= t_obs)

## §7 — Driscoll-Kraay sensitivity
```

```r
for (L in 1:3) {
  m_dk <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean + 
sovciss_lag1 +",
                                 ctrl, "| country + year")),
                data = df_lag, panel.id = ~country + year, vcov = DK(L))

}

## §8 — MDE on gamma_4

mde_from <- function(m, param, power = 0.80, df_t = 10) {
  se <- sqrt(vcov(m)[param, param])
  (qt(0.975, df_t) + qt(power, df_t)) * se
}

mde_g4 <- mde_from(h4_joint, "sovciss_lag1")
gamma4 <- coef(h4_joint)["sovciss_lag1"]


## ==========================================================================
## 5. H5V2.R — H5: non-linearity of the stress-expenditure relation
##    Quadratic model, turning point, marginal effects, spline, leave-
one-out, MDE
## ==========================================================================

## Packages
library(readxl); library(dplyr); library(tidyr)
library(fixest); library(car)

## §1 — Import via header patterns

path <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données 
groupées.xlsx"
df_raw <- read_excel(path, sheet = "Données")
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]

h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))

pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf(
    "Motif '%s' -> %d colonne(s) [%s]. Corrige ce motif.",
    pattern, length(idx), paste(names(df_raw)[idx], collapse = " | ")))
  idx
}

map <- c(
  country          = "^country$",
  year             = "^time$|^year$",
  cult_exp_gdp     = "cultural services exp \\(%gdp\\)",
  cult_exp_totexp  = "cult serv exp \\(%tot",
  debt_gdp         = "debt stock",
  gdp_growth       = "growth rate",
  unemployment     = "unemployment",
  sovciss_mean     = "sovciss.*annual",
  wgi_ge           = "wgi",
```

```r
  eurozone         = "^eurozone$"
)

sel <- sapply(map, pick)
df  <- df_raw[, sel]; names(df) <- names(map)

df <- df %>%
  mutate(country = trimws(country)) %>%
  fill(country, .direction = "down") %>%
  filter(!is.na(year)) %>%
  mutate(country = as.factor(country),
         year     = as.integer(year),
         eurozone = as.integer(eurozone)) %>%
  filter(eurozone == 1) %>%
  mutate(country = droplevels(country))

stopifnot(nrow(df) == 209, length(unique(df$country)) == 11)
sd_by_country <- tapply(df$sovciss_mean, df$country, sd, na.rm = TRUE)
stopifnot(all(sd_by_country > 0.001))

## §2 — Centred terms

mu_sov <- mean(df$sovciss_mean)
df <- df %>% mutate(
  sovciss_c  = sovciss_mean - mu_sov,
  sovciss_c2 = sovciss_c^2
)

corr_lin_quad <- cor(df$sovciss_c, df$sovciss_c2)

ctrl <- "gdp_growth + debt_gdp + unemployment"

## §3 — Nested models: linear vs quadratic

m_lin <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c +", ctrl,
                                "| country + year")),
               data = df, cluster = ~country)

m_quad <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c + sovciss_c2
+", ctrl,
                                 "| country + year")),
                data = df, cluster = ~country)

print(etable(m_lin, m_quad,
             headers = c("Lineaire (rappel H1)", "H5 : quadratique"),
             tex = FALSE))

print(wald(m_quad, keep = "sovciss_c2"))

## §4 — Turning point

tp <- deltaMethod(m_quad, "-sovciss_c/(2*sovciss_c2)", vcov =
vcov(m_quad))
print(tp)
tp_raw <- mu_sov + tp$Estimate
```

```r
## §5 — Marginal effects across the stress distribution (Table 6b)

qs   <- c(.10, .50, .75, .90)
pcts <- quantile(df$sovciss_mean, qs)
lab  <- c("10th", "50th (median)", "75th", "90th")

est <- se <- tv <- pv <- numeric(length(qs))
for (i in seq_along(qs)) {
  cp     <- pcts[[i]] - mu_sov
  me     <- deltaMethod(m_quad,
                        sprintf("sovciss_c + 2*%.6f*sovciss_c2", cp),
                        vcov = vcov(m_quad))
  est[i] <- me$Estimate; se[i] <- me$SE
  tv[i]  <- est[i] / se[i]
  pv[i]  <- 2 * pt(-abs(tv[i]), df = 10)
}
star <- ifelse(pv < .01, "***", ifelse(pv < .05, "**", ifelse(pv < .1,
"*", "")))

me6 <- data.frame(
  "Percentile of SovCISS" = lab,
  "SovCISS"               = sprintf("%.3f", pcts),
  "Marginal effect"       = sprintf("%+.3f%s", est, star),
  "SE"                    = sprintf("%.3f", se),
  "t"                     = sprintf("%+.2f", tv),
  "p-value"               = sprintf("%.3f", pv),
  check.names = FALSE, stringsAsFactors = FALSE)
me6[] <- lapply(me6, function(x) gsub("-", "\u2212", x, fixed = TRUE))

## §6 — Driscoll-Kraay sensitivity

m_quad_dk <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c +
sovciss_c2 +", ctrl,
                                    "| country + year")),
                   data = df, panel.id = ~country + year, vcov = DK(2))

for (L in 1:3) {
  m_L <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c + sovciss_c2
+", ctrl,
                                "| country + year")),
               data = df, panel.id = ~country + year, vcov = DK(L))

}

## §7 — Enumerated wild cluster bootstrap on gamma_5

m_r <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_c +", ctrl,
                              "| country + year")),
             data = df)
df$fitted_r <- fitted(m_r); df$resid_r <- resid(m_r)

t_obs <- coeftable(m_quad)["sovciss_c2", "t value"]

cl <- sort(unique(as.character(df$country)))
W  <- as.matrix(expand.grid(rep(list(c(-1, 1)), length(cl))))
```

```r
fml_boot <- as.formula(paste("y_star ~ sovciss_c + sovciss_c2 +", ctrl,
                             "| country + year"))

t_boot <- apply(W, 1, function(w) {
  names(w) <- cl
  df$y_star <- df$fitted_r + df$resid_r * w[as.character(df$country)]
  coeftable(feols(fml_boot, data = df, cluster = ~country))["sovciss_c2",
"t value"]
})

p_boot_h5 <- mean(abs(t_boot) >= abs(t_obs))

## §8 — Threshold spline (knots at q75 and q90)

q75 <- quantile(df$sovciss_mean, 0.75)
q90 <- quantile(df$sovciss_mean, 0.90)
df <- df %>% mutate(
  kink75 = pmax(sovciss_mean - q75, 0),
  kink90 = pmax(sovciss_mean - q90, 0)
)

m_spl75 <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean + kink75
+", ctrl,
                                  "| country + year")),
                 data = df, cluster = ~country)
m_spl90 <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean + kink90
+", ctrl,
                                  "| country + year")),
                 data = df, cluster = ~country)

print(etable(m_quad, m_spl75, m_spl90,
             headers = c("Quadratique", "Spline (nœud q75)", "Spline
(nœud q90)"),
             tex = FALSE))

print(deltaMethod(m_spl75, "sovciss_mean + kink75", vcov =
vcov(m_spl75)))

print(deltaMethod(m_spl90, "sovciss_mean + kink90", vcov =
vcov(m_spl90)))

## §9 — Leave-one-country-out on gamma_5

gamma5_full <- coef(m_quad)["sovciss_c2"]

loo_g5 <- sapply(levels(df$country), function(cc) {
  d <- filter(df, country != cc)
  d <- d %>% mutate(sc = sovciss_mean - mean(sovciss_mean),
                    sc2 = sc^2)
  m <- feols(as.formula(paste("cult_exp_gdp ~ sc + sc2 +", ctrl,
                              "| country + year")),
             data = d, cluster = ~country)
  coef(m)["sc2"]
})

print(sort(round(loo_g5, 3)))
```

```r
## §10 — Leverage diagnostic: observations above the 90th percentile

top <- df %>% filter(sovciss_mean > q90) %>%
  select(country, year, sovciss_mean, cult_exp_gdp) %>%
  arrange(desc(sovciss_mean))
print(as.data.frame(top))

print(table(droplevels(top$country)))

top_by_country <- table(droplevels(top$country))
top_sorted <- sort(top_by_country, decreasing = TRUE)
top_label <- paste(
  paste(names(top_sorted)[1:3], top_sorted[1:3], sep = " "),
  collapse = ", ")
other_n <- sum(top_sorted[-(1:3)])
top_decile_str <- sprintf("%s, other %d", top_label, other_n)

## §11 — Alternative dependent variable (% of total expenditure)

m_quad_tot <- feols(as.formula(paste("cult_exp_totexp ~ sovciss_c + 
sovciss_c2 +",
                                     ctrl, "| country + year")),
                    data = df, cluster = ~country)

print(etable(m_quad, m_quad_tot,
             headers = c("H5 : %PIB", "H5 : %depenses totales"),
             tex = FALSE))

## §12 — MDE on gamma_5

mde_from <- function(m, param, power = 0.80, df_t = 10) {
  se <- sqrt(vcov(m)[param, param])
  (qt(0.975, df_t) + qt(power, df_t)) * se
}

mde_g5 <- mde_from(m_quad, "sovciss_c2")


## =====================================================================
## 6. RobustesseFull.R — Robustness battery R1-R8 and encompassing test
##    Placebo lead, alternative stress measures, jackknife, encompassing, 
per-country slopes
## =====================================================================

## Packages
library(readxl); library(dplyr); library(tidyr)
library(fixest)

## §1 — Import via header patterns

path <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données 
groupées.xlsx"
df_raw <- read_excel(path, sheet = "Données")
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]
```

```r
pick <- function(pattern, required = TRUE) {
  hits <- grep(pattern, names(df_raw), ignore.case = TRUE, perl = TRUE,
value = TRUE)
  if (length(hits) == 0L && !required) return(NA_character_)
  if (length(hits) != 1L)
    stop("Motif '", pattern, "' -> ", length(hits), " colonne(s) : ",
         paste(hits, collapse = " | "), ". Corrige ce motif.")
  hits
}

col_country  <- pick("^country$")
col_year     <- pick("^time$|^year$")
col_cegdp    <- pick("(?s)cultural services exp.*%\\s*gdp")
col_ctot_r   <- pick("(?s)cult serv exp.*%\\s*tot")
col_growth   <- pick("^growth rate")
col_debt     <- pick("^debt stock")
col_unemp    <- pick("^unemployment")
col_sov      <- pick("^sovciss\\s*\\(annual")
col_sovmax   <- pick("^sovciss\\s*\\(max")
col_te       <- pick("^cultural services expenditure\\s*\\(m€\\)\\s*$")

col_ciss_c   <- pick("^\\s*newciss\\s*\\(annual\\)")
col_ciss_ea  <- pick("newciss.*euro")
col_ideo     <- pick("parlgov|gov.*ideo|ideo.*gov|left.*right")
col_hicp     <- pick("^hicp$")
col_pop      <- pick("^pop$")

col_crisis08    <- pick("2008 crisis")
col_crisisdebt  <- pick("^debt crisis")
col_crisiscov   <- pick("covid crisis")
col_ez <- pick("^\\s*eurozone\\s*$")

df <- df_raw |>
  transmute(
    country         = .data[[col_country]],
    year            = suppressWarnings(as.integer(.data[[col_year]])),
    cult_exp_gdp    = suppressWarnings(as.numeric(.data[[col_cegdp]])),
    cult_exp_totexp = suppressWarnings(as.numeric(.data[[col_ctot_r]])),
    cult_te         = suppressWarnings(as.numeric(.data[[col_te]])),
    gdp_growth      = suppressWarnings(as.numeric(.data[[col_growth]])),
    debt_gdp        = suppressWarnings(as.numeric(.data[[col_debt]])),
    unemployment    = suppressWarnings(as.numeric(.data[[col_unemp]])),
    sovciss_mean    = suppressWarnings(as.numeric(.data[[col_sov]])),
    sovciss_max     = suppressWarnings(as.numeric(.data[[col_sovmax]])),
    ciss_country    = suppressWarnings(as.numeric(.data[[col_ciss_c]])),
    ciss_ea         = suppressWarnings(as.numeric(.data[[col_ciss_ea]])),
    gov_lr          = suppressWarnings(as.numeric(.data[[col_ideo]])),
    hicp            = suppressWarnings(as.numeric(.data[[col_hicp]])),
    pop             = suppressWarnings(as.numeric(.data[[col_pop]])),
    eurozone        = suppressWarnings(as.integer(.data[[col_ez]])),
    crisis_2008     =
suppressWarnings(as.numeric(.data[[col_crisis08]])),
    crisis_debt     =
suppressWarnings(as.numeric(.data[[col_crisisdebt]])),
    crisis_covid    =
suppressWarnings(as.numeric(.data[[col_crisiscov]]))
```

```r
  ) |>
  mutate(country = trimws(country)) |>
  fill(country, .direction = "down") |>
  filter(!is.na(year), year >= 2006L, year <= 2024L) |>
  filter(eurozone == 1L) |>
  mutate(country = droplevels(as.factor(country)))

stopifnot(nrow(df) == 209, n_distinct(df$country) == 11)

ctrl <- "gdp_growth + debt_gdp + unemployment"

## §2 — Derived variables (lead for the placebo; real per-capita terms
for R5)

df <- df |>
  arrange(country, year) |>
  group_by(country) |>
  mutate(sovciss_lead1 = dplyr::lead(sovciss_mean, 1)) |>
  ungroup()
stopifnot(sum(is.na(df$sovciss_lead1)) == 11)

df <- df |>
  mutate(
    nom_gdp      = cult_te / (cult_exp_gdp / 100),
    real_cult_pc = cult_te * 100 / hicp / pop,
    real_gdp_pc  = nom_gdp * 100 / hicp / pop,
    log_cult_pc  = log(real_cult_pc),
    log_gdp_pc   = log(real_gdp_pc)
  )

## §3 — Reference baseline (direct check against Table 3a)

f_base_h1 <- as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl, "|
country + year"))
f_base_h2 <- as.formula(paste("cult_exp_totexp ~ sovciss_mean +", ctrl,
"| country + year"))

b_h1 <- feols(f_base_h1, df, cluster = ~country)
b_h2 <- feols(f_base_h2, df, cluster = ~country)

print(etable(b_h1, b_h2, headers = c("H1 baseline", "H2 baseline")))

## §4 — Variance decomposition of SovCISS

fe_dec <- lm(sovciss_mean ~ factor(country) + factor(year), data = df)
r2_fe  <- summary(fe_dec)$r.squared
sd_res <- sd(resid(fe_dec))

## §5 — Placebo: one lead of SovCISS

df_lead <- filter(df, !is.na(sovciss_lead1))
stopifnot(nrow(df_lead) == 198)

m_placebo <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +
sovciss_lead1 +",
                                    ctrl, "| country + year")),
```

```r
                        data = df_lead, cluster = ~country)

print(etable(m_placebo, headers = "Placebo : SovCISS(t) + SovCISS(t+1)"))

## §6 — R1: country-level CISS (10 countries, Greece excluded)

df10 <- filter(df, !is.na(ciss_country))
stopifnot(n_distinct(df10$country) == 10)

f_ciss10_h1 <- as.formula(paste("cult_exp_gdp ~ ciss_country +", ctrl, "|
country + year"))
f_ciss10_h2 <- as.formula(paste("cult_exp_totexp ~ ciss_country +", ctrl,
"| country + year"))

r1_sov10_h1 <- feols(f_base_h1, df10, cluster = ~country)
r1_ciss10_h1 <- feols(f_ciss10_h1, df10, cluster = ~country)
r1_sov10_h2 <- feols(f_base_h2, df10, cluster = ~country)
r1_ciss10_h2 <- feols(f_ciss10_h2, df10, cluster = ~country)

print(etable(b_h1, r1_sov10_h1, r1_ciss10_h1,
             headers = c("Baseline (11c, SovCISS)", "Pont (10c,
SovCISS)", "R1: CISS (10c)")))

print(etable(b_h2, r1_sov10_h2, r1_ciss10_h2,
             headers = c("Baseline (11c, SovCISS)", "Pont (10c,
SovCISS)", "R1: CISS (10c)")))

## §7 — R2: euro-area CISS (country fixed effects only)

f_sv_noyfe_h1 <- as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
"| country"))
f_ea_h1       <- as.formula(paste("cult_exp_gdp ~ ciss_ea +", ctrl, "|
country"))
f_sv_noyfe_h2 <- as.formula(paste("cult_exp_totexp ~ sovciss_mean +",
ctrl, "| country"))
f_ea_h2       <- as.formula(paste("cult_exp_totexp ~ ciss_ea +", ctrl, "|
country"))

r2_sv_h1 <- feols(f_sv_noyfe_h1, df, cluster = ~country)
r2_ea_h1 <- feols(f_ea_h1, df, cluster = ~country)
r2_sv_h2 <- feols(f_sv_noyfe_h2, df, cluster = ~country)
r2_ea_h2 <- feols(f_ea_h2, df, cluster = ~country)

print(etable(r2_sv_h1, r2_ea_h1, headers = c("SovCISS (sans FE annee)",
"R2: CISS EA")))

print(etable(r2_sv_h2, r2_ea_h2, headers = c("SovCISS (sans FE annee)",
"R2: CISS EA")))

## §8 — R3: annual maximum of SovCISS

f_max_h1 <- as.formula(paste("cult_exp_gdp ~ sovciss_max +", ctrl, "|
country + year"))
f_max_h2 <- as.formula(paste("cult_exp_totexp ~ sovciss_max +", ctrl, "|
country + year"))
```

```r
r3_max_h1 <- feols(f_max_h1, df, cluster = ~country)
r3_max_h2 <- feols(f_max_h2, df, cluster = ~country)

print(etable(b_h1, r3_max_h1, b_h2, r3_max_h2,
             headers = c("H1 base", "H1 max", "H2 base", "H2 max")))

## §9 — R4: government-ideology control (ParlGov)

f_id_h1 <- as.formula(paste("cult_exp_gdp ~ sovciss_mean + gov_lr +",
ctrl, "| country + year"))
f_id_h2 <- as.formula(paste("cult_exp_totexp ~ sovciss_mean + gov_lr +",
ctrl, "| country + year"))

r4_id_h1 <- feols(f_id_h1, df, cluster = ~country)
r4_id_h2 <- feols(f_id_h2, df, cluster = ~country)

print(etable(b_h1, r4_id_h1, b_h2, r4_id_h2,
             headers = c("H1 base", "H1 + ideo", "H2 base", "H2 +
ideo")))

## §10 — R5: log real cultural expenditure per capita (H1 only)

f_lg_h1 <- as.formula(paste("log_cult_pc ~ sovciss_mean + log_gdp_pc +",
ctrl, "| country + year"))
r5_lg_h1 <- feols(f_lg_h1, df, cluster = ~country)

print(etable(b_h1, r5_lg_h1, headers = c("H1 base (part de PIB)", "R5:
log reel p.c.")))

delta_sov <- median(tapply(df$sovciss_mean, df$country,
                           function(x) diff(range(x, na.rm = TRUE))))
eff_base  <- 100 * coef(b_h1)["sovciss_mean"] * delta_sov /
mean(df$cult_exp_gdp, na.rm = TRUE)
eff_r5    <- 100 * (exp(coef(r5_lg_h1)["sovciss_mean"] * delta_sov) - 1)

## MDE for R5 (log real per capita)
se_r5     <- se(r5_lg_h1)["sovciss_mean"]

t_crit    <- qt(0.95, df = 10)
t_power   <- qt(0.80, df = 10)

mde_semi  <- (t_crit + t_power) * se_r5

delta_sov <- median(tapply(df$sovciss_mean,
                           df$country,
                           function(x) diff(range(x, na.rm = TRUE))))

mde_pct   <- (exp(mde_semi * delta_sov) - 1) * 100

## §11 — R6: leave-one-country-out (H1 and H2)

countries <- sort(unique(as.character(df$country)))

loo_h1 <- sapply(countries, function(cc)
  coef(feols(f_base_h1, filter(df, country != cc), cluster = ~country))
["sovciss_mean"])
```

```r
loo_h2 <- sapply(countries, function(cc)
  coef(feols(f_base_h2, filter(df, country != cc), cluster = ~country))
["sovciss_mean"])

print(round(sort(loo_h1), 4))

print(round(sort(loo_h2), 4))

## §12 — R7: excluding COVID years (2020-2021)

df_nocov <- filter(df, !year %in% 2020:2021)

r7_nc_h1 <- feols(f_base_h1, df_nocov, cluster = ~country)
r7_nc_h2 <- feols(f_base_h2, df_nocov, cluster = ~country)

print(etable(b_h1, r7_nc_h1, b_h2, r7_nc_h2,
             headers = c("H1 base", "H1 sans COVID", "H2 base", "H2 sans
COVID")))

## §13 — R8: two-way clustering (country x year)

r8_tw_h1 <- feols(f_base_h1, df, cluster = ~country + year)
r8_tw_h2 <- feols(f_base_h2, df, cluster = ~country + year)

print(etable(b_h1, r8_tw_h1, b_h2, r8_tw_h2,
             headers = c("H1 pays", "H1 bidirectionnel", "H2 pays", "H2
bidirectionnel")))

## §14 — Encompassing: crisis dummies vs SovCISS (H1)

within_year_var <- df |> group_by(year) |>
  summarise(across(c(crisis_2008, crisis_debt, crisis_covid), ~var(.x,
na.rm = TRUE)))
pure_time_dummies <- all(colSums(within_year_var[, -1], na.rm = TRUE) <
1e-10)
fe_bench <- if (pure_time_dummies) "| country" else "| country + year"

m_dum  <- feols(as.formula(paste("cult_exp_gdp ~ crisis_2008 +
crisis_debt + crisis_covid +",
                                 ctrl, fe_bench)),
                data = df, cluster = ~country)
m_ciss <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
fe_bench)),
                data = df, cluster = ~country)
m_both <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +
crisis_2008 + crisis_debt + crisis_covid +",
                                 ctrl, fe_bench)),
                data = df, cluster = ~country)

print(etable(m_dum, m_ciss, m_both,
             headers = c("Dummies (replication)", "SovCISS", "Les deux
(encompassing)"),
             fitstat = ~ n + r2 + wr2 + aic + bic))

print(wald(m_both, keep = "crisis"))
```

```r
print(wald(m_both, keep = "sovciss"))

## §14bis — Per-country bivariate slopes on both measures; Greek
contraction 2008-2010

stopifnot("log_cult_pc" %in% names(df))

pente <- function(d, yvar) {
  d <- d[!is.na(d[[yvar]]) & !is.na(d$sovciss_mean), ]
  if (nrow(d) < 3) return(NA_real_)
  unname(coef(lm(d[[yvar]] ~ d$sovciss_mean))[2])
}

slopes <- do.call(rbind, lapply(split(df, df$country), function(d) {
  data.frame(country    = as.character(d$country[1]),
             slope_gdp  = pente(d, "cult_exp_gdp"),
             slope_real = pente(d, "log_cult_pc"),
             stringsAsFactors = FALSE)
}))
slopes <- slopes[order(slopes$slope_real), ]
row.names(slopes) <- NULL

print(slopes, digits = 3)

gr <- df[df$country == "Greece" & df$year %in% c(2008, 2010),
         c("year", "cult_te")]
print(gr)


## ====================================================================
## 7. tableau_7c.R — Placebo: leads of SovCISS (strict-exogeneity check)
##    Joint Wald test and enumerated bootstrap on the two leads
## ====================================================================

library(readxl); library(dplyr); library(tidyr)
library(fixest)

data_path  <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données
groupées.xlsx"
sheet_name <- "Données"

df_raw <- read_excel(data_path, sheet = sheet_name)
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]
h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))

pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf(
    "Pattern '%s' -> %d column(s) [%s].",
    pattern, length(idx), paste(names(df_raw)[idx], collapse = " | ")))
  idx
}

map <- c(
  country       = "^country$",
  year          = "^time$|^year$",
```

```r
  cult_exp_gdp = "cultural services exp \\(%gdp\\)",
  debt_gdp     = "debt stock",
  gdp_growth   = "growth rate",
  unemployment = "unemployment",
  sovciss_mean = "sovciss.*annual",
  eurozone     = "^eurozone$"
)

df <- df_raw[, sapply(map, pick)]; names(df) <- names(map)
df <- df %>%
  mutate(country = trimws(country)) %>%
  fill(country, .direction = "down") %>%
  filter(!is.na(year)) %>%
  mutate(country = as.factor(country), year = as.integer(year), eurozone
= as.integer(eurozone)) %>%
  filter(eurozone == 1) %>%
  mutate(country = droplevels(country))
stopifnot(nrow(df) == 209, length(unique(df$country)) == 11)

ctrl <- "gdp_growth + debt_gdp + unemployment"

## Leads t+1 and t+2 by country
df <- df %>%
  arrange(country, year) %>%
  group_by(country) %>%
  mutate(sov_f1 = dplyr::lead(sovciss_mean, 1),
         sov_f2 = dplyr::lead(sovciss_mean, 2)) %>%
  ungroup()

d1 <- filter(df, !is.na(sov_f1))
d2 <- filter(df, !is.na(sov_f1) & !is.na(sov_f2))
stopifnot(nrow(d1) == 198, nrow(d2) == 187)

## Models
m_lead1 <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean + sov_f1
+", ctrl,
                                  "| country + year")), data = d1,
cluster = ~country)

m_lead2 <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean + sov_f1 +
sov_f2 +", ctrl,
                                  "| country + year")), data = d2,
cluster = ~country)

print(coeftable(m_lead1)["sov_f1", ])
print(coeftable(m_lead2)[c("sov_f1","sov_f2"), ])

## Joint Wald F on the two leads (clustered vcov, df = G-1 = 10)
joint_F <- function(m, params) {
  b <- coef(m)[params]; V <- vcov(m)[params, params]; q <- length(params)
  as.numeric(t(b) %*% solve(V) %*% b) / q
}
params <- c("sov_f1", "sov_f2")
Fobs   <- joint_F(m_lead2, params)
p_ana  <- pf(Fobs, length(params), 10, lower.tail = FALSE)
```

```r
## Enumerated wild cluster bootstrap on the joint null (both leads = 0)
m_r <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
                              "| country + year")), data = d2)
d2$fitted_r <- fitted(m_r); d2$resid_r <- resid(m_r)

cl <- sort(unique(as.character(d2$country)))
W  <- as.matrix(expand.grid(rep(list(c(-1, 1)), length(cl))))
fml_b <- as.formula(paste("y_star ~ sovciss_mean + sov_f1 + sov_f2 +",
ctrl,
                          "| country + year"))

F_boot <- apply(W, 1, function(w) {
  names(w) <- cl
  d2$y_star <- d2$fitted_r + d2$resid_r * w[as.character(d2$country)]
  joint_F(feols(fml_b, data = d2, cluster = ~country), params)
})
p_boot <- mean(F_boot >= Fobs)


## =====================================================================
## 8. within_vs_diffe_rences.R — Within vs first-difference estimator
(Section 3.7)
## =====================================================================

library(readxl); library(dplyr); library(tidyr); library(fixest)

data_path  <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données
groupées.xlsx"
sheet_name <- "Données"
df_raw <- read_excel(data_path, sheet = sheet_name)
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]
h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))
pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf("Pattern '%s' -> %d col(s).",
pattern, length(idx)))
  idx }
map <- c(country = "^country$", year = "^time$|^year$",
         cult_exp_gdp = "cultural services exp \\(%gdp\\)",
         debt_gdp = "debt stock", gdp_growth = "growth rate",
         unemployment = "unemployment", sovciss_mean = "sovciss.*annual",
         eurozone = "^eurozone$")
sel <- sapply(map, pick); df <- df_raw[, sel]; names(df) <- names(map)
df <- df %>%
  mutate(country = trimws(country)) %>% fill(country, .direction =
"down") %>%
  filter(!is.na(year)) %>%
  mutate(country = as.factor(country), year = as.integer(year),
         eurozone = as.integer(eurozone)) %>%
  filter(eurozone == 1) %>% mutate(country = droplevels(factor(country)))
stopifnot(nrow(df) == 209, length(unique(df$country)) == 11)
ctrl <- "gdp_growth + debt_gdp + unemployment"

## Within estimator
m_fe <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
                               "| country + year")),
```

```r
                 data = df, cluster = ~country)

## First differences (GDP growth kept in level, already a rate)
df_fd <- df %>% arrange(country, year) %>% group_by(country) %>%
  mutate(d_cult  = cult_exp_gdp - dplyr::lag(cult_exp_gdp),
         d_sov   = sovciss_mean - dplyr::lag(sovciss_mean),
         d_debt  = debt_gdp     - dplyr::lag(debt_gdp),
         d_unemp = unemployment - dplyr::lag(unemployment)) %>%
  ungroup() %>% filter(!is.na(d_sov))

m_fd <- feols(d_cult ~ d_sov + gdp_growth + d_debt + d_unemp | year,
                 data = df_fd, cluster = ~country)

print(etable(m_fe, m_fd, headers = c("Within", "First diff."),
                keep = c("sovciss_mean", "d_sov")))

b_fe <- coef(m_fe)["sovciss_mean"]; b_fd <- coef(m_fd)["d_sov"]


## ======================================================================
## 9. tests_specification.R — Specification tests
##    CIPS, poolability, Mean Group and CCEMG, Nickell bias, Anderson-
Hsiao IV
## ======================================================================

library(plm)

data_path  <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données 
groupées.xlsx"
sheet_name <- "Données"

## 6.1 — Import

df_raw <- read_excel(data_path, sheet = sheet_name)
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]
h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))

pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf(
    "Pattern '%s' -> %d column(s) [%s].",
    pattern, length(idx), paste(names(df_raw)[idx], collapse = " | ")))
  idx
}

map <- c(
  country         = "^country$",
  year            = "^time$|^year$",
  cult_exp_gdp    = "cultural services exp \\(%gdp\\)",
  cult_exp_totexp = "cult serv exp \\(%tot",
  debt_gdp        = "debt stock",
  gdp_growth      = "growth rate",
  unemployment    = "unemployment",
  sovciss_mean    = "sovciss.*annual",
  eurozone        = "^eurozone$"
)
```

```r
sel <- sapply(map, pick)
df  <- df_raw[, sel]; names(df) <- names(map)

df <- df %>%
  mutate(country = trimws(country)) %>%
  fill(country, .direction = "down") %>%
  filter(!is.na(year)) %>%
  mutate(country = as.factor(country),
         year     = as.integer(year),
         eurozone = as.integer(eurozone)) %>%
  filter(eurozone == 1) %>%
  mutate(country = droplevels(country))
stopifnot(nrow(df) == 209, length(unique(df$country)) == 11)

pdat <- pdata.frame(df, index = c("country", "year"))
f_spec <- cult_exp_gdp ~ sovciss_mean + gdp_growth + debt_gdp +
unemployment

## 6.2 — CIPS panel unit root tests (Pesaran, 2007)

vars_ur <- c("cult_exp_gdp", "cult_exp_totexp", "sovciss_mean")

cips_safe <- function(x, type, lags, label) {
  res <- tryCatch(cipstest(x, lags = lags, type = type, model = "cmg"),
                  error   = function(e) NULL,
                  warning = function(w) suppressWarnings(
                    cipstest(x, lags = lags, type = type, model =
"cmg")))
  invisible(res)
}

for (v in vars_ur) {
  for (tp in c("drift", "trend")) cips_safe(pdat[[v]], tp, 1, v)
}

for (v in vars_ur) {
  cips_safe(diff(pdat[[v]]), "drift", 1, paste0("d.", v))
}

## 6.3 — Poolability: Chow-type F test of slope homogeneity

pool_res <- tryCatch(pooltest(f_spec, data = pdat, model = "within"),
                     error = function(e) {NULL})
if (!is.null(pool_res)) print(pool_res)

## 6.4 — Mean Group (Pesaran & Smith, 1995) and CCEMG (Pesaran, 2006)

m_mg <- tryCatch(pmg(f_spec, data = pdat, model = "mg"),
                 error = function(e) {NULL})
if (!is.null(m_mg)) { print(summary(m_mg)) }

m_ccemg <- tryCatch(pmg(f_spec, data = pdat, model = "cmg"),
                    error = function(e) {NULL})
if (!is.null(m_ccemg)) { print(summary(m_ccemg)) }
```

```r
if (!is.null(m_mg)) {
  ind <- m_mg$indcoef
  if ("sovciss_mean" %in% rownames(ind)) {
    sl <- sort(ind["sovciss_mean", ])

    print(round(sl, 4))

  }
}

## 6.5 — Estimator comparison (Table 9)

get_row <- function(m, label, param = "sovciss_mean") {
  if (is.null(m)) return(NULL)
  ct <- summary(m)$CoefTable
  if (!param %in% rownames(ct)) return(NULL)
  data.frame(Estimator = label,
             Coefficient = sprintf("%+.4f", ct[param, 1]),
             SE          = sprintf("%.4f",  ct[param, 2]),
             `p-value`   = sprintf("%.3f",  ct[param, 4]),
             check.names = FALSE, stringsAsFactors = FALSE)
}

fe_ct  <- coeftable(feols(cult_exp_gdp ~ sovciss_mean + gdp_growth +
debt_gdp +
                            unemployment | country + year,
                          data = df, cluster = ~country))
fe_row <- data.frame(
  Estimator   = "Two-way FE (baseline)",
  Coefficient = sprintf("%+.4f", fe_ct["sovciss_mean", 1]),
  SE          = sprintf("%.4f",  fe_ct["sovciss_mean", 2]),
  `p-value`   = sprintf("%.3f",  fe_ct["sovciss_mean", 4]),
  check.names = FALSE, stringsAsFactors = FALSE)

cmp <- do.call(rbind, Filter(Negate(is.null), list(
  fe_row,
  get_row(m_mg,    "Mean Group"),
  get_row(m_ccemg, "CCEMG (common factors)")
)))
cmp[] <- lapply(cmp, function(x) gsub("-", "\u2212", x, fixed = TRUE))

print(cmp)

## 6.6 — Persistence, Nickell bias, and infeasibility of GMM

df_ar <- df %>% arrange(country, year) %>% group_by(country) %>%
  mutate(cult_lag1 = dplyr::lag(cult_exp_gdp, 1)) %>% ungroup() %>%
  filter(!is.na(cult_lag1))

m_ar <- feols(cult_exp_gdp ~ cult_lag1 + sovciss_mean + gdp_growth +
debt_gdp +
                unemployment | country + year,
              data = df_ar, cluster = ~country)
print(etable(m_ar, headers = "With lagged dependent variable (biased)"))

rho <- coef(m_ar)["cult_lag1"]
```

```r
T_len <- length(unique(df$year))
nickell <- -(1 + rho) / (T_len - 1)

## First-difference specification

df_fd <- df %>% arrange(country, year) %>% group_by(country) %>%
  mutate(d_cult   = cult_exp_gdp - dplyr::lag(cult_exp_gdp),
         d_sov    = sovciss_mean - dplyr::lag(sovciss_mean),
         d_debt   = debt_gdp     - dplyr::lag(debt_gdp),
         d_unemp  = unemployment - dplyr::lag(unemployment)) %>%
  ungroup() %>% filter(!is.na(d_sov))

m_fd <- feols(d_cult ~ d_sov + gdp_growth + d_debt + d_unemp | year,
              data = df_fd, cluster = ~country)
etable(m_fd, headers = "First differences")

## Anderson-Hsiao IV: first differences, D.SovCISS instrumented by
SovCISS(t-2)

df_ah <- df %>% arrange(country, year) %>% group_by(country) %>%
  mutate(d_cult    = cult_exp_gdp - dplyr::lag(cult_exp_gdp),
         d_sov_iv  = sovciss_mean - dplyr::lag(sovciss_mean),
         d_debt    = debt_gdp     - dplyr::lag(debt_gdp),
         d_unemp   = unemployment - dplyr::lag(unemployment),
         sov_l2    = dplyr::lag(sovciss_mean, 2)) %>%
  ungroup() %>% filter(!is.na(d_sov_iv), !is.na(sov_l2))

m_ah <- feols(d_cult ~ gdp_growth + d_debt + d_unemp |
                year | d_sov_iv ~ sov_l2,
              data = df_ah, cluster = ~country)

fs      <- fitstat(m_ah, c("ivf", "wh"))
ah_fs_F <- as.numeric(fs$`ivf1::d_sov_iv`$stat)

chow_F <- if (!is.null(pool_res)) pool_res$statistic else NA
chow_p <- if (!is.null(pool_res)) pool_res$p.value else NA

extra_rows <- data.frame(
  Estimator = c("MG slope dispersion (SovCISS)", "Chow poolability (F)",
                "First differences (β, d_sov)",
                "Anderson-Hsiao IV (β)",
                "GMM instrument count"),
  Coefficient = c(
    sprintf("range [%+.4f; %+.4f]", min(sl), max(sl)),
    sprintf("F=%.2f", chow_F),
    sprintf("%+.4f", coef(m_fd)["d_sov"]),
    sprintf("%+.4f", coef(m_ah)["fit_d_sov_iv"]),
    sprintf("%d (N=%d)", (T_len-2)*(T_len-1)/2,
length(unique(df$country)))
  ),
  SE = c(sprintf("sd=%.4f", sd(sl)), "",
         sprintf("%.4f", se(m_fd)["d_sov"]),
         sprintf("%.4f", se(m_ah)["fit_d_sov_iv"]),
         ""),
  `p-value`  = c(sprintf("%d/%d neg.", sum(sl<0), length(sl)),
                sprintf("%.3f", chow_p),
```

```r
                sprintf("%.3f", pvalue(m_fd)["d_sov"]),
                sprintf("first-stage F=%.1f", ah_fs_F),
                ""),
  check.names = FALSE, stringsAsFactors = FALSE
)
cmp_full <- rbind(cmp, extra_rows)


## =========================================================================
## 10. panel_e_tendu.R — Extended 16-country panel
##    Slope homogeneity, Chow tests, CCEMG, and every hypothesis re-
estimated
## =========================================================================

library(readxl); library(dplyr); library(tidyr)
library(fixest); library(plm); library(car)

data_path  <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données
groupées.xlsx"
sheet_name <- "Données"

## 7.1 — Import and group indicator

df_raw <- read_excel(data_path, sheet = sheet_name)
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]
h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))

pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf(
    "Pattern '%s' -> %d column(s) [%s].",
    pattern, length(idx), paste(names(df_raw)[idx], collapse = " | ")))
  idx
}

map <- c(
  country          = "^country$",
  year             = "^time$|^year$",
  cult_exp_gdp     = "cultural services exp \\(%gdp\\)",
  cult_exp_totexp  = "cult serv exp \\(%tot",
  debt_gdp         = "debt stock",
  gdp_growth       = "growth rate",
  unemployment     = "unemployment",
  sovciss_mean     = "sovciss.*annual",
  eurozone         = "^eurozone$"
)

sel <- sapply(map, pick)
dx  <- df_raw[, sel]; names(dx) <- names(map)

dx <- dx %>%
  mutate(country = trimws(country)) %>%
  fill(country, .direction = "down") %>%
  filter(!is.na(year)) %>%
  mutate(country  = as.factor(country),
         year     = as.integer(year),
```

```r
        eurozone = as.integer(eurozone))

stopifnot(nrow(dx) == 304, length(unique(dx$country)) == 16)
stopifnot(sum(!is.na(dx$unemployment)) == 304)

n_ez  <- n_distinct(dx$country[dx$eurozone == 1])
n_nez <- n_distinct(dx$country[dx$eurozone == 0])

stopifnot(n_ez == 11, n_nez == 5)

ctrl <- "gdp_growth + debt_gdp + unemployment"

## 7.2 — Descriptive comparison (Table 10a)

desc_grp <- dx %>%
  mutate(Group = ifelse(eurozone == 1, "Eurozone (11)", "Non-eurozone
(5)")) %>%
  group_by(Group) %>%
  summarise(N               = n(),
            `Cult. exp (% GDP)` = mean(cult_exp_gdp, na.rm = TRUE),
            `SovCISS (mean)`    = mean(sovciss_mean, na.rm = TRUE),
            `SovCISS (max)`     = max(sovciss_mean,  na.rm = TRUE),
            `Debt (% GDP)`      = mean(debt_gdp,     na.rm = TRUE),
            `GDP growth (%)`    = mean(gdp_growth,   na.rm = TRUE),
            `Unemployment (%)`  = mean(unemployment, na.rm = TRUE),
            .groups = "drop") %>%
  mutate(across(where(is.numeric), ~round(., 3)))

print(as.data.frame(desc_grp))

## 7.3 — Interaction test: SovCISS x Eurozone

m_ext_base <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +",
ctrl,
                                     "| country + year")),
                    data = dx, cluster = ~country)

m_ext_int  <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +
sovciss_mean:eurozone +",
                                     ctrl, "| country + year")),
                    data = dx, cluster = ~country)

print(etable(m_ext_base, m_ext_int,
             headers = c("Extended pooled", "+ SovCISS x Eurozone")))

print(deltaMethod(m_ext_int, "sovciss_mean + `sovciss_mean:eurozone`",
                  vcov = vcov(m_ext_int)))

## 7.4 — Chow test: all coefficients interacted

m_chow <- feols(as.formula(paste(
  "cult_exp_gdp ~ sovciss_mean + gdp_growth + debt_gdp + unemployment +",
  "sovciss_mean:eurozone + gdp_growth:eurozone + debt_gdp:eurozone +",
  "unemployment:eurozone | country + year")),
  data = dx, cluster = ~country)
```

```r
print(etable(m_chow, headers = "Fully interacted"))
print(wald(m_chow, keep = ":eurozone"))

## 7.5 — Separate estimation by group

m_ez  <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
                                "| country + year")),
               data = filter(dx, eurozone == 1), cluster = ~country)
m_nez <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl,
                                "| country + year")),
               data = filter(dx, eurozone == 0), cluster = ~country)

print(etable(m_ez, m_nez, m_ext_base,
             headers = c("Eurozone only (11)", "Non-eurozone only (5)",
"Pooled (16)")))

## 7.6 — Mean Group slopes by group

pdx <- pdata.frame(dx, index = c("country", "year"))
f_spec_ext <- cult_exp_gdp ~ sovciss_mean + gdp_growth + debt_gdp +
unemployment

m_mg_ext <- tryCatch(pmg(f_spec_ext, data = pdx, model = "mg"),
                     error = function(e) {NULL})

if (!is.null(m_mg_ext)) {

  print(summary(m_mg_ext))

  sl  <- m_mg_ext$indcoef["sovciss_mean", ]
  grp <- dx %>% distinct(country, eurozone)
  ez_flag <- grp$eurozone[match(names(sl), as.character(grp$country))]

  out <- data.frame(country = names(sl),
                    slope   = round(as.numeric(sl), 4),
                    group   = ifelse(ez_flag == 1, "EZ", "non-EZ"))
  print(out[order(out$slope), ], row.names = FALSE)

  d  <- mean(sl[ez_flag == 1]) - mean(sl[ez_flag == 0])
  se <- sqrt(var(sl[ez_flag == 1]) / sum(ez_flag == 1) +
             var(sl[ez_flag == 0]) / sum(ez_flag == 0))

}

## 7.7 — CCEMG on the extended sample: shock diffusion

m_ccemg_ext <- tryCatch(pmg(f_spec_ext, data = pdx, model = "cmg"),
                        error = function(e) {NULL})
if (!is.null(m_ccemg_ext)) {

  print(summary(m_ccemg_ext))
}

m_plm_ext <- plm(f_spec_ext, data = pdx, model = "within", effect =
"twoways")
```

```r
print(pcdtest(m_plm_ext, test = "cd"))

## 7.8 — Wild cluster bootstrap on the interaction (Webb weights)

boot_webb <- function(data, dep, rhs_full, rhs_restr, param, B = 9999,
seed = 2026) {
  set.seed(seed)
  fe  <- "| country + year"
  m_r <- feols(as.formula(paste(dep, "~", rhs_restr, fe)), data = data)
  data$fitted_r <- fitted(m_r); data$resid_r <- resid(m_r)

  m_f   <- feols(as.formula(paste(dep, "~", rhs_full, fe)), data = data,
cluster = ~country)
  t_obs <- coeftable(m_f)[param, "t value"]

  cl <- sort(unique(as.character(data$country))); G <- length(cl)
  w6 <- c(-sqrt(1.5), -1, -sqrt(0.5), sqrt(0.5), 1, sqrt(1.5))
  fml_b <- as.formula(paste("y_star ~", rhs_full, fe))

  t_boot <- replicate(B, {
    w <- sample(w6, G, replace = TRUE); names(w) <- cl
    data$y_star <- data$fitted_r + data$resid_r *
w[as.character(data$country)]
    coeftable(feols(fml_b, data = data, cluster = ~country))[param, "t
value"]
  })
  list(t_obs = t_obs, p_two = mean(abs(t_boot) >= abs(t_obs)))
}

b_int_ext <- boot_webb(
  dx, "cult_exp_gdp",
  paste("sovciss_mean + sovciss_mean:eurozone +", ctrl),
  paste("sovciss_mean +", ctrl),
  "sovciss_mean:eurozone")

## Additional rows for Table 10b (Chow test and MG slope dispersion)
chow_w <- wald(m_chow, keep = ":eurozone", print = FALSE)
sl16   <- m_mg_ext$indcoef["sovciss_mean", ]

extra_10b <- data.frame(
  term = c("Chow F (partition, 4 df)", "  p (partition)",
           "MG slope range (16 countries)"),
  `Pooled (16)`       = c(sprintf("%.2f", chow_w$stat), sprintf("%.3f",
chow_w$p),
                         sprintf("[%+.3f; %+.3f]", min(sl16),
max(sl16))),
  `+ Interaction`     = c("", "", ""),
  `Fully interacted` = c("", "", ""),
  `Eurozone (11)`     = c("", "", ""),
  `Non-eurozone (5)` = c("", "", ""),
  check.names = FALSE, stringsAsFactors = FALSE
)
attr(extra_10b, "position") <- NULL

## PART 8 — All hypotheses re-estimated on the extended sample (Table
10d)
```

```r
library(dplyr); library(tidyr); library(fixest)
library(car)

data_path  <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données groupées.xlsx"
sheet_name <- "Données"

stopifnot(exists("dx"))
ctrl <- "gdp_growth + debt_gdp + unemployment"

## 8.1 — Derived variables on the extended panel

h <- tolower(gsub("\\s+", " ", trimws(names(df_raw))))
pick <- function(pattern) {
  idx <- which(grepl(pattern, h))
  if (length(idx) != 1) stop(sprintf("Pattern '%s' -> %d column(s).", pattern, length(idx)))
  idx
}
extra <- df_raw[, sapply(c(country = "^country$", year = "^time$|^year$",
                           cult_meur = "expenditure \\(m€\\)$",
                           hicp = "^hicp$", pop = "^pop$"), pick)]
names(extra) <- c("country", "year", "cult_meur", "hicp", "pop")
extra <- extra %>% mutate(country = trimws(country)) %>%
  fill(country, .direction = "down") %>%
  filter(!is.na(year)) %>%
  mutate(country = as.factor(country), year = as.integer(year))

dx <- dx %>% left_join(extra, by = c("country", "year"))

mu_ext <- mean(dx$sovciss_mean, na.rm = TRUE)

dx <- dx %>%
  arrange(country, year) %>%
  group_by(country) %>%
  mutate(sovciss_lag1 = dplyr::lag(sovciss_mean, 1)) %>%
  ungroup() %>%
  mutate(
    sc   = sovciss_mean - mu_ext,
    sc2 = sc^2,
    nom_gdp      = cult_meur / (cult_exp_gdp / 100),
    log_cult_pc = log(cult_meur * 100 / hicp / pop),
    log_gdp_pc   = log(nom_gdp  * 100 / hicp / pop)
  )

dx_lag <- filter(dx, !is.na(sovciss_lag1))
stopifnot(nrow(dx_lag) == 304 - 16)

## 8.2 — H1, H2, H4, H5 and R5 on 16 countries

e_h1 <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +", ctrl, "| country + year")),
              data = dx, cluster = ~country)
```

```r
e_h2 <- feols(as.formula(paste("cult_exp_totexp ~ sovciss_mean +", ctrl,
"| country + year")),
              data = dx, cluster = ~country)

e_h4 <- feols(as.formula(paste("cult_exp_gdp ~ sovciss_mean +
sovciss_lag1 +", ctrl,
                              "| country + year")),
              data = dx_lag, cluster = ~country)

e_h5 <- feols(as.formula(paste("cult_exp_gdp ~ sc + sc2 +", ctrl, "|
country + year")),
              data = dx, cluster = ~country)

e_r5 <- feols(as.formula(paste("log_cult_pc ~ sovciss_mean + log_gdp_pc
+", ctrl,
                              "| country + year")),
              data = dx, cluster = ~country)

print(etable(e_h1, e_h2, e_h4, e_h5, e_r5,
             headers = c("H1 (16c)", "H2 (16c)", "H4 (16c)", "H5 (16c)",
"R5 (16c)")))

print(deltaMethod(e_h4, "sovciss_mean + sovciss_lag1", vcov =
vcov(e_h4)))
print(wald(e_h4, keep = "sovciss"))

print(wald(e_h5, keep = "sc2"))

## 8.3 — Table 10d: side-by-side comparison

grab <- function(m, param) {
  ct <- coeftable(m)
  c(est = ct[param, 1], se = ct[param, 2], p = ct[param, 4])
}

fmt <- function(v) {
  star <- ifelse(v["p"] < .01, "***", ifelse(v["p"] < .05, "**",
                                      ifelse(v["p"] < .1, "*",
"")))
  sprintf("%+.4f%s (%.4f)", v["est"], star, v["se"])
}

base11 <- list(
  H1 = c(est = -0.0249, se = 0.0615, p = 0.695),
  H2 = c(est = -0.0365, se = 0.1096, p = 0.746),
  H4 = c(est = -0.0261, se = 0.0457, p = 0.581),
  H5 = c(est = -0.0795, se = 0.1169, p = 0.512),
  R5 = c(est = -0.1875, se = 0.1823, p = 0.328)
)

ext16 <- list(
  H1 = grab(e_h1, "sovciss_mean"),
  H2 = grab(e_h2, "sovciss_mean"),
  H4 = grab(e_h4, "sovciss_lag1"),
  H5 = grab(e_h5, "sc2"),
  R5 = grab(e_r5, "sovciss_mean")
```

```r
)

ccemg16 <- grab(m_ccemg_ext, "sovciss_mean")

cmp10c <- data.frame(
  Specification = c(
    "H1  Cultural exp. (% GDP), SovCISS",
    "H2  Cultural exp. (% total exp.), SovCISS",
    "H4  Lagged SovCISS (t−1)",
    "H5  Quadratic term SovCISS squared",
    "R5  log real expenditure per capita",
    "CCEMG (common factors), SovCISS"),
  `Eurozone (11)`  = c(sapply(base11, fmt),
                        fmt(c(est = -0.0215, se = 0.0377, p = 0.568))),
  `Extended (16)`  = c(sapply(ext16, fmt), fmt(ccemg16)),
  N_11 = c(209, 209, 198, 209, 209, 209),
  N_16 = c(nobs(e_h1), nobs(e_h2), nobs(e_h4), nobs(e_h5), nobs(e_r5), 
304),
  check.names = FALSE, stringsAsFactors = FALSE, row.names = NULL)
cmp10c[] <- lapply(cmp10c, function(x) gsub("-", "\u2212", 
as.character(x), fixed = TRUE))

print(cmp10c, row.names = FALSE)


## =====================================================================
## 11. tableau_10c.R — Kapetanios (2003) sequential poolability of the 
SovCISS slope
##    Panel A: 11 euro-area countries; Panel B: 16 EU countries
## =====================================================================

library(readxl); library(dplyr); library(tidyr); library(fixest)

path <- "/Users/charlesmaes/Desktop/UCL/Mémoire/Données/Données 
groupées.xlsx"

df_raw <- read_excel(path, sheet = "Données")
df_raw <- df_raw[, !grepl("^\\.\\.\\.", names(df_raw))]

pick <- function(pattern) {
  hits <- grep(pattern, names(df_raw), ignore.case = TRUE, perl = TRUE, 
value = TRUE)
  if (length(hits) != 1L)
    stop("Motif '", pattern, "' -> ", length(hits), " colonne(s) : ",
         paste(hits, collapse = " | "))
  hits
}

cols <- c(country = "^country$", year = "^time$|^year$",
          cult_exp_gdp = "(?s)cultural services exp.*%\\s*gdp",
          gdp_growth = "^growth rate", debt_gdp = "^debt stock",
          unemployment = "^unemployment", sovciss_mean = "^sovciss\\s*\\
(annual",
          eurozone = "^\\s*eurozone\\s*$")

sel <- vapply(cols, pick, character(1))
```

```r
d <- df_raw[, sel]; names(d) <- names(cols)

d <- d |>
  mutate(country = trimws(country)) |>
  fill(country, .direction = "down") |>
  filter(!is.na(year)) |>
  mutate(across(c(year, eurozone), as.integer),
         across(c(cult_exp_gdp, gdp_growth, debt_gdp, unemployment,
sovciss_mean),
                ~ suppressWarnings(as.numeric(.x))),
         country = as.factor(country)) |>
  filter(year >= 2006L, year <= 2024L)

d16 <- d |> mutate(country = droplevels(country))
d11 <- d |> filter(eurozone == 1L) |> mutate(country =
droplevels(country))

stopifnot(nrow(d16) == 304, n_distinct(d16$country) == 16,
          nrow(d11) == 209, n_distinct(d11$country) == 11)

ctrl <- "gdp_growth + debt_gdp + unemployment"

## Joint Wald test on the country-specific SovCISS slopes

sov_wald <- function(data, active, min_size = 3) {
  if (length(active) < min_size) return(list(F = NA_real_, p = NA_real_))
  dd <- droplevels(subset(data, country %in% active))
  f  <- as.formula(paste("cult_exp_gdp ~ sovciss_mean +
sovciss_mean:country +",
                         ctrl, "| country + year"))
  m <- tryCatch(feols(f, data = dd, vcov = "hetero"), error = function(e)
NULL)
  if (is.null(m)) return(list(F = NA_real_, p = NA_real_))
  w <- tryCatch(wald(m, keep = "sovciss_mean:country"), error =
function(e) NULL)
  if (is.null(w)) return(list(F = NA_real_, p = NA_real_))
  list(F = unname(w[["stat"]]), p = unname(w[["p"]]))
}

kapetanios <- function(data, alpha = 0.05, min_size = 3) {
  active <- levels(droplevels(data$country))
  hist <- data.frame(step = integer(), dropped = character(),
                     F = numeric(), p = numeric(), stringsAsFactors =
FALSE)
  step <- 0
  repeat {
    step <- step + 1
    res <- sov_wald(data, active, min_size)
    if (is.na(res$p) || res$p >= alpha || length(active) <= min_size) {
      return(list(active = active, history = hist,
                  final_F = res$F, final_p = res$p, n_final =
length(active)))
    }
    contrib <- vapply(active, function(cc)
      sov_wald(data, setdiff(active, cc), min_size)$F, numeric(1))
    drop <- names(which.min(contrib))
```

```r
    hist <- rbind(hist, data.frame(step = step, dropped = drop,
                                   F = res$F, p = res$p))
    active <- setdiff(active, drop)
  }
}

res11 <- kapetanios(d11)
res16 <- kapetanios(d16)

## Assemble Table 10c

fmt_p <- function(p) ifelse(p < 1e-4, "< 0.0001", sprintf("%.4f", p))

rows <- function(r, label) {
  h <- r$history
  body <- data.frame(
    Step       = as.character(h$step),
    Excluded   = h$dropped,
    `F stat`   = sprintf("%.3f", h$F),
    `p value`  = fmt_p(h$p),
    check.names = FALSE
  )
  stop_row <- data.frame(
    Step = "stop", Excluded = "— (null not rejected)",
    `F stat` = sprintf("%.3f", r$final_F), `p value` = fmt_p(r$final_p),
    check.names = FALSE
  )
  rbind(body, stop_row)
}

A <- rows(res11, "A"); B <- rows(res16, "B")
tab <- rbind(A, B)

ez <- d16 |> distinct(country, eurozone) |> mutate(country =
as.character(country))
grp16 <- ez |> filter(country %in% res16$active)
inside  <- sort(grp16$country[grp16$eurozone == 1L])
outside <- sort(grp16$country[grp16$eurozone == 0L])

note <- paste0(
  "Sequential exclusion following Kapetanios. At each step the joint Wald
test ",
  "of equality of the country-specific SovCISS slopes is computed; the
country ",
  "whose removal lowers the statistic most is excluded.
Heteroskedasticity-robust ",
  "variance (clustering by country is infeasible with one interaction per
cluster). ",
  "Controls: GDP growth, public debt, unemployment; country and year
fixed effects. ",
  "Panel B poolable group (", res16$n_final, " countries). Euro area: ",
  paste(inside, collapse = ", "), "; outside: ", paste(outside, collapse
= ", "),
  ". The procedure does not correct for multiple testing."
)
```

```r
print(tab, row.names = FALSE)


## ============================================================================
## 12. Monte_Carlo.R — Monte Carlo: finite-sample properties of the two-
way FE estimator
##    N = 11, T = 19; bias and size under strict exogeneity and under 
feedback
## ============================================================================

library(fixest); library(dplyr)

## 9.1 — One simulated panel

sim_one <- function(N = 11, T = 19, beta = 0,
                    rho_x = 0.85, rho_e = 0.795,
                    feedback = 0, sd_e = 0.30, sd_x = 0.20) {

  mu  <- rnorm(N, 0, 1.0)
  lam <- rnorm(T, 0, 0.5)

  x <- e <- matrix(0, nrow = N, ncol = T)

  x[, 1] <- rnorm(N, 0, sd_x / sqrt(1 - rho_x^2))
  e[, 1] <- rnorm(N, 0, sd_e / sqrt(1 - rho_e^2))

  for (t in 2:T) {
    e[, t] <- rho_e * e[, t - 1] + rnorm(N, 0, sd_e)
    x[, t] <- rho_x * x[, t - 1] + feedback * e[, t - 1] + rnorm(N, 0, 
sd_x)
  }

  y <- outer(mu, rep(1, T)) + outer(rep(1, N), lam) + beta * x + e

  d <- data.frame(
    country = factor(rep(1:N, times = T)),
    year    = factor(rep(1:T, each  = N)),
    y       = as.vector(y),
    x       = as.vector(x)
  )

  m <- feols(y ~ x | country + year, data = d, cluster = ~country)
  ct <- coeftable(m)
  c(est = ct["x", 1], se = ct["x", 2], p = ct["x", 4])
}

## 9.2 — Run both DGPs (A: strictly exogenous; B: feedback from the 
lagged error)

run_mc <- function(R = 2000, feedback, label, beta = 0, seed = 2026) {
  set.seed(seed)
  out <- t(replicate(R, sim_one(beta = beta, feedback = feedback)))
  bias    <- mean(out[, "est"]) - beta
  sd_est  <- sd(out[, "est"])
  mean_se <- mean(out[, "se"])
  size    <- mean(out[, "p"] < 0.05)
```

```r
  invisible(list(bias = bias, sd = sd_est, se = mean_se, size = size))
}

mc_a <- run_mc(R = 2000, feedback = 0.0,
               label = "DGP A: x strictly exogenous")

mc_b <- run_mc(R = 2000, feedback = 0.5,
               label = "DGP B: x predetermined (feedback from lagged
error)")

## 9.3 — Bias as N grows, then as T grows (Nickell-type bias is O(1/T))

for (Nn in c(11, 25, 50, 100)) {
  set.seed(2026)
  out <- t(replicate(500, sim_one(N = Nn, feedback = 0.5)))
}

for (Tt in c(19, 40, 80, 160)) {
  set.seed(2026)
  out <- t(replicate(500, sim_one(T = Tt, feedback = 0.5)))
}
```