# The Very Long-Run Costs of Financial Crises*

Karsten Müller† Chenzi Xu‡ Mohamed Lehbib§

March 2026

**Abstract**

We provide new empirical evidence of the long-run effects of financial crises on the real economy by examining output dynamics over unprecedented time horizons. To do so, we introduce the Global Macroeconomic Database, which covers 243 countries and extends back to the origins of modern data collection. Using this extended historical coverage, we find that banking crises are associated with statistically significant contractions in real GDP that persist for five decades following the initial shock. These estimated effects are considerably larger than those found in previous studies, which have typically been constrained by shorter time horizons and less comprehensive data. We find a similarly prolonged disruption in both human and physical capital accumulation while fiscal revenues deteriorate, spending increases, and debt accumulates. These disruptions are particularly pronounced during globally synchronized crises and in emerging economies. The effects are consistent with the view that recoveries are slower when alternative balance sheets cannot absorb financial losses, suggesting that financial constraints are a central mechanism. The magnitudes of these persistent output losses imply that the true economic costs of financial instability have been substantially underestimated in the literature.

*Astrid Tonnet provided excellent research assistance.

†National University of Singapore, NUS Business School and Risk Management Institute, & CEPR, kmueller@nus.edu.sg

‡University of California, Berkeley & NBER & CEPR; chenzixu@berkeley.edu

§National University of Singapore, NUS Business School, lehbib@nus.edu.sg

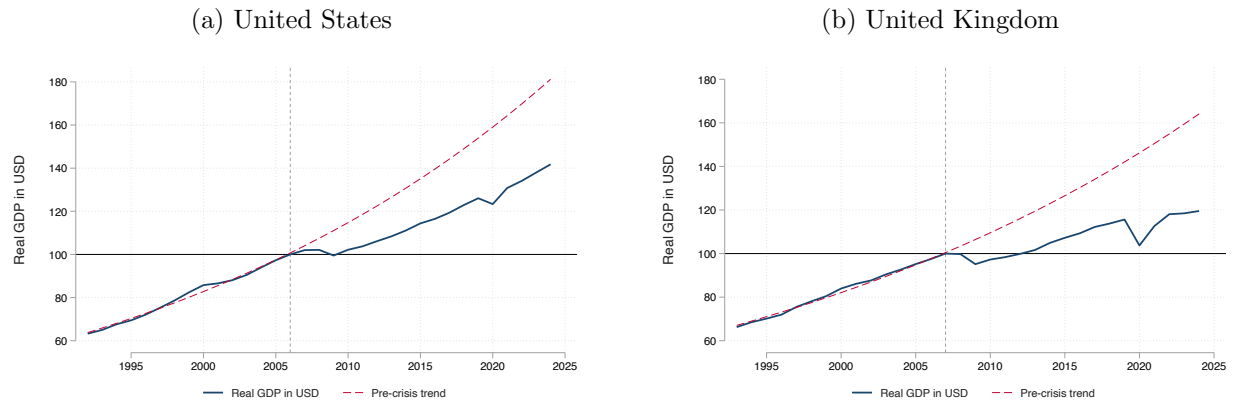
# 1 Introduction

The global financial crisis of 2008 is widely viewed as a defining macroeconomic event of the modern era. In many economies around the world, output fell sharply during the crisis, and the subsequent recovery was slow. Almost two decades later, economic activity remains below the trajectory implied by pre-crisis growth trends (Ball, 2014).

Figure 1 illustrates this pattern for two widely discussed cases, the United States and the United Kingdom. In each figure, we plot the realized path of real GDP (solid blue line) along with the projected path based on the growth trend pre-crisis (dashed red line). In both countries, output fell sharply during the crisis and recovered only gradually. In addition, the recovery in both cases has been incomplete as real GDP remains substantially below the path implied by pre-crisis growth.

These patterns have revived the debate on whether financial crises generate permanent economic damage (e.g., Cerra and Saxena, 2008; Xu, 2022) or whether economies eventually return to levels implied by their previous growth paths (e.g., Davis and Weinstein, 2002). More broadly, how persistent are the macroeconomic consequences of financial crises in general? Answering this question requires data that combine two features that are rarely available simultaneously: broad cross-country coverage and sufficiently long time series that trace economic dynamics over horizons measured in decades rather than years.

Figure 1: Output Dynamics Following the Global Financial Crisis




*Notes:* This figure illustrates the evolution of real GDP around the 2007-2008 Global Financial crisis for the United States and United Kingdom, respectively. Each panel plots real GDP (solid line) normalized to 100 in the year preceding the crisis, together with the counterfactual trajectory implied by the pre-crisis trend (dashed line). The pre-crisis trend is estimated using average growth between years $t-15$ and $t-1$, where $t$ denotes the crisis year. Vertical dashed lines indicate the onset of the banking crisis. Crisis dates follow the banking crisis chronology constructed in this paper.

This paper provides such evidence using a newly constructed Global Macro Database (GMD) that links modern national accounts to historical macroeconomic series for the full set of countries for which macroeconomic data exist. The database covers 243 countries and extends back to the origins of modern statistical measurement, combining the cross-country breadth of modern macroe-

conomic datasets with the temporal depth of historical reconstructions. The resulting sample is substantially larger than those of prior work (Appendix Figure A.1 provides a detailed comparison).

Studying the impact of crises over this extended sample also requires a banking crisis chronology with comparable breadth. Each existing chronology in the literature is designed with specific time periods or country sets in mind, and none provides the coverage necessary to match the full span of the GMD. To address this, we construct a new harmonized chronology by taking the union of the four principal crisis datasets used in the literature—Reinhart and Rogoff (2009*b*), Laeven and Valencia (2013), Jordà, Schularick and Taylor (2017), Baron, Verner and Xiong (2021), along with our own additions based on primary source research. Each source applies different definitions and dating conventions so we impose a uniform crisis definition and revisit each candidate episode to reconcile discrepancies across datasets. In addition, we adjudicate disagreements across sources by systematically reviewing the historical evidence. The resulting chronology is deliberately inclusive in its candidate episodes but requires consistent historical documentation to be classified as a crisis.

We estimate the dynamic effects of banking crises using local projections, which allow us to trace the evolution of output over long horizons without imposing strong parametric restrictions on adjustment dynamics. A central feature of our design is to include flexible controls of pre-crisis growth dynamics over short, medium, and long horizons, which capture the economic build-ups that often precede crises, and which ensure that post-crisis estimates are not conflated with the unwinding of pre-existing trends. Implementing this design with the required lag structure while still obtaining precise estimates at horizons of up to fifty years demands large samples, which the breadth and length of the GMD is uniquely positioned to provide.

Our central finding is that banking crises generate extremely persistent output losses. Real GDP declines sharply following a crisis and remains significantly below its pre-crisis trajectory for decades. Even fifty years after the initial shock, average output remains substantially depressed relative to the counterfactual path implied by pre-crisis growth dynamics. We also find that there is no evidence of differential pre-trends in output dynamics in the twenty years prior to crisis outset, which helps to lend credibility to interpreting these effects as causal. These long-run losses are considerably larger than those estimated in previous studies that examine shorter horizons or rely on more limited samples.

The results are robust to a wide range of specification checks. First, the effects are stable across lag structures ranging from five to twenty-five years of pre-crisis growth history. Second, the results are similar when we constrain the data to a balanced panel that requires a country to be present in all the years of a given estimation. Third, conditioning on contemporaneous wars, currency crises, and sovereign debt crises leaves the impulse responses essentially unchanged, as does the inclusion of a broad set of macroeconomic controls including inflation, interest rates,

credit aggregates, investment, and government debt. Fourth, excluding crises associated with the two World Wars and the Great Depression does not attenuate the findings. Finally, the results are robust to using each of the four underlying crisis chronologies individually, as well as their union, in place of our new harmonized indicator.

We also subject the results to a range of alternative econometric approaches. A placebo exercise that preserves the overall frequency of crises but randomizes their occurrence generates impulse responses centered near zero at all horizons, indicating that the estimated losses are not an artifact of underlying output persistence or slow-moving macroeconomic trends. The results are also robust to estimating local projections in levels and long differences rather than cumulative growth rates, and to estimating the dynamic response using the Pesaran and Smith (1995) mean group estimator, which allows crisis effects to vary freely across countries.

Having established the magnitude and robustness of these long-run output losses, in the last section of the paper, we turn to a set of analyses that shed light on the mechanisms underlying their persistence. Our approach is not to isolate a single structural channel, but rather to document a set of empirical patterns that collectively point to the role of financial constraints that interrupt capital accumulation and increase debt burdens in generating long-run economic scarring.

We begin by examining the dynamics of crises themselves. We find that financial crises are followed by elevated probabilities of subsequent crises, consistent with the view that crises can weaken financial systems and increase future vulnerability. However, these dynamics do not account for the persistence of output losses: even conditioning on the full path of crises, the long-run effects remain large. This suggests that persistence reflects more than repeated crisis realizations, and instead points to lasting damage from the initial shock.

We next turn to capital accumulation and technology adoption. We find that financial crises are followed by sustained declines in both physical and human capital investment as well as a broad index of technology adoption, indicating prolonged disruptions to the processes that underpin long-run growth. Investment rates remain depressed for extended periods, and measures of human capital accumulation exhibit similarly persistent shortfalls. These patterns are consistent with a tightening of financial constraints that limits the ability of firms and households to undertake forward-looking investments.

These costs are evident on government balance sheets as well. On average, government revenues as a share of GDP declines persistently, which combines with a short-term increase in government spending to generate deficits and debt accumulation. While revenue shares remain persistently lower, expenditures and debt burdens both eventually decline, suggesting that on average, countries reduce the size of the public sector following these disruptions.

In the final part of the paper, we examine heterogeneity in the persistence of crisis effects along

dimensions that capture the availability of alternative balance sheets and the capacity of economies to absorb financial shocks. We first distinguish between crises that occur in isolation and those that are part of broader episodes of global financial distress, as well as between crises occurring in core financial centers (such as the United Kingdom in the prewar era and the United States in the postwar period) and those in peripheral economies. Persistence is substantially more pronounced in globally synchronized crises and in episodes centered in major financial hubs, consistent with the idea that disruptions to global financial intermediation propagate shocks more widely and hinder recovery.

We then compare advanced and emerging economies, where differences in fiscal capacity and financial development shape the ability to stabilize the economy and repair private sector balance sheets. Emerging economies exhibit more persistent output losses on average, consistent with more limited policy space and weaker mechanisms for absorbing financial distress.

Taken together, these findings point to a common interpretation: financial crises generate persistent output losses by impairing balance sheets and tightening financial constraints, thereby disrupting capital accumulation and slowing recovery. As a result, the long-run costs of crises are substantially larger than what is implied by short-horizon analyses, underscoring the importance of policies aimed at crisis prevention and rapid balance sheet repair.

*Related Literature.*

This paper relates to a large literature documenting the macroeconomic consequences of financial crises. A central finding in this literature is that banking crises are associated with substantial declines in economic activity. Using long-run historical data, Reinhart and Rogoff (2009*b*); Barro (2006); Barro and Ursúa (2008) document that financial crises and other rare disasters coincide with sharp contractions in output and consumption. More recent work using panel variation across countries shows that credit booms often precede financial crises and that these crises are followed by prolonged recessions and slow recoveries (Schularick and Taylor, 2012; Jordà, Schularick and Taylor, 2013). Similarly, Cerra and Saxena (2008) show that output losses following financial crises are highly persistent relative to other shocks such as wars.

A related literature examines the broader consequences of financial crises across a range of economic and political outcomes. Crises have been shown to generate persistent effects on potential output and demand dynamics (Blanchard, Cerutti and Summers, 2015; Cerra, Fatás and Saxena, 2023), as well as on international trade patterns (Xu, 2022). Beyond real economic activity, financial crises also have important political consequences: Funke, Schularick and Trebesch (2016) document that they are followed by sharp increases in political extremism, while Guiso, Herrera, Morelli and Sonno (2017); Quincy and Xu (2026) show that financial and economic distress contributes to rising support for populist movements. While these contributions establish that financial crises have

wide-ranging effects, with the exception of Xu (2022), the empirical evidence has largely focused on horizons of up to a decade and feature little discussion of the mechanisms for persistence.

This paper also relates to a large body of work emphasizing the role of financial frictions in amplifying macroeconomic fluctuations. Models of the financial accelerator and collateral constraints show how disruptions in credit markets can lead to large declines in real activity (Bernanke and Gertler, 1989; Kiyotaki and Moore, 1997). Consistent with these mechanisms, empirical evidence shows that financial conditions, measured for instance with credit spreads, strongly amplify business cycle fluctuations, particularly during downturns (Krishnamurthy and Muir, 2025). Other work highlights the importance of the composition of credit in shaping macroeconomic outcomes. For example, expansions in household credit and lending to financially constrained sectors are associated with more severe downturns and a higher likelihood of financial crises (Mian, Sufi and Verner, 2017; Müller and Verner, 2024). While this literature provides a compelling account of how financial disruptions generate large short- and medium-run effects, it offers more limited guidance on why the effects of financial crises remain so persistent over multi-decade horizons.

Finally, our analysis relates to a broader literature on recovery from large macroeconomic shocks. A growing body of work emphasizes that post-crisis dynamics depend critically on the ability of economies to stabilize and rebuild after the initial contraction (e.g., Romer and Romer, 2019). Our findings complement this perspective by showing that the persistence is systematically related to fiscal capacity and the broader functioning of financial intermediation.

## 2 The Global Macro Database

This paper introduces the Global Macro Database (GMD), a new harmonized dataset designed to support empirical analysis of macroeconomic dynamics over very long horizons. The database combines modern national accounts with historical macroeconomic series to construct consistent annual measures of output and related aggregates for a large cross-section of countries over extended time spans. Its central objective is to make credible inference about macroeconomic shocks—particularly rare but consequential events such as financial crises—possible over decadal horizons.

The approach in constructing the GMD is motivated by a fundamental limitation of existing macroeconomic datasets. Most widely used sources either offer broad cross-country coverage over relatively short post-Bretton Woods samples, or long historical series for a narrow set of economies. This limitation is compounded by the substantial practical challenges involved in combining historical and modern data sources, which differ in units of account, price deflators, territorial definitions, and statistical methodologies, making the construction of consistent long-run cross-country series nontrivial. As a result, empirical work on the long-run consequences of macroeconomic disruptions has typically relied on small samples, truncated horizons, or highly selective country coverage.

The GMD is designed to overcome these constraints by unifying over one hundred historical and contemporary sources into a single, harmonized database suitable for long-horizon analysis. Coverage varies across countries and time, reflecting the availability of historical sources, but the database is explicitly constructed to maximize both temporal depth and cross-sectional breadth. This feature is central to the analysis that follows.

**Coverage and scope.** The Global Macro Database provides annual macroeconomic data for 243 countries and territories, with coverage extending as far back as the medieval period for some variables and countries, and continuing through the present. The core focus of the database is real economic activity, measured primarily through real GDP and closely related aggregates, which form the backbone of the analysis in this paper.

In addition, the database includes a broad set of complementary macroeconomic variables, including prices, population, public finances, labor markets, and external accounts, which support the study of mechanisms and longer-run adjustment dynamics.

**Conceptual design.** The GMD features three key design elements: long-run continuity, cross-country comparability, and replication transparency.

First, the database is designed to link all available data about a country into a unified time series whenever possible. This continuity is essential for tracing the persistent effects of shocks whose consequences may unfold slowly over time. Second, we harmonize all series across countries to allow for cross-sectional comparison. Our database provides values in nominal and real values, as well as local currency and USD. Third, the database is designed as a durable research infrastructure rather than a one-off compilation. All series are constructed using documented sources and open source reproducible procedures, with versioning and revision protocols intended to preserve historical consistency as new data become available.

## 2.1 Data construction

Constructing a unified database from 118 sources requires addressing substantial heterogeneity in coverage, units, and methodological approaches. Our construction methodology prioritizes preserving the growth dynamics within each source while anchoring levels to the most reliable contemporary benchmarks. This approach preserves growth rate variation over absolute historical levels, which are measured with considerably greater noise in earlier periods, and which avoids artificial breaks that would arise from simply concatenating disparate sources.

We implement this principle through a systematic ratio-splicing procedure. For each country and variable, we establish a benchmark series from the highest-quality modern source—typically the IMF World Economic Outlook or national statistical offices for the post-1990 period. We then

extend this series backwards by sequentially linking earlier historical sources, which preserves their growth rates while adjusting their levels to ensure continuity at splice points.

Formally, consider a benchmark series $y_t^{(a)}$ and an earlier historical series $y_t^{(b)}$ that overlap in year $\tau$. We construct the spliced series as:

$$\tilde{y}_t^{(b)} = y_t^{(b)} \times \frac{y_\tau^{(a)}}{y_\tau^{(b)}} \quad \text{for } t < \tau$$

The final series uses $\tilde{y}_t^{(b)}$ for $t < \tau$ and $y_t^{(a)}$ for $t \geq \tau$. This procedure maintains the growth rates from the historical source—which typically reflect careful work by economic historians—while eliminating level discontinuities that would otherwise arise from differences in base years, currency units, or statistical methodology.

When multiple historical sources are available for a given period, we apply this procedure sequentially, working backwards from the benchmark and prioritizing sources with broader temporal coverage and more comprehensive documentation. In the small number of cases in which adjacent sources do not overlap, we link them using median growth rates calculated from surrounding years. These instances are explicitly flagged in our documentation, and we provide ensitivity analyses using alternative linking procedures.

To ensure full transparency and replicability, we provide a comprehensive Technical Appendix that documents the construction of each series and reports the underlying source data alongside the final GMD series for every country and variable. In addition, all code used to construct the database is publicly available. This structure allows researchers to trace each series back to its original sources, evaluate the impact of alternative splicing or harmonization choices, and reproduce or modify the database for their own purposes. The GMD is therefore designed not only as a dataset for the present analysis, but as a flexible and extensible research infrastructure for future empirical work on long-run macroeconomic dynamics.

**Data revision policy.** Contemporary macroeconomic data are subject to continuous revision as statistical agencies refine their estimates, which poses a challenge for long-run databases that splice historical and modern sources. Unrestricted propagation of these revisions backward through the entire time series would undermine replicability and may render earlier empirical findings incomparable to those using updated data.

We address this challenge by fixing historical values while allowing contemporary data to incorporate ongoing revisions. Specifically, we establish 2019 as a reference year: values prior to 2019 remain fixed unless a clear error is discovered or demonstrably superior historical evidence becomes available, while values from 2020 onward update with each new release from contemporary sources.

This approach ensures that the historical portions of our series—which form the foundation for identifying long-run effects—remain stable, while recent data incorporate the latest available information. Routine national accounts updates thus affect only the most recent years and do not cascade through earlier periods.

In addition to this updating process, all versions of the database are archived and publicly accessible, allowing researchers to replicate earlier results or assess the sensitivity of findings to data revisions. Comprehensive documentation of source selection, splicing procedures, and any revisions is provided in the Technical Appendix.

**Variable definitions.** Beyond mechanical splicing, constructing consistent long-run series requires addressing conceptual differences across sources—for example, in the treatment of government output, subsistence production, or territorial coverage. Where feasible, we harmonize historical series to align with modern national accounts concepts as defined by the System of National Accounts 2008. When full harmonization is not possible—typically for pre-1950 data in which original source definitions may diverge substantially from modern conventions—we preserve the original definition while documenting the discrepancy and ensuring internal consistency within each country's time series.

All real variables are expressed in constant 2015 local currency units. For historical periods in which only nominal series exist, we apply deflators constructed from the closest available price index, prioritizing those from the same underlying source to maintain consistency. Population series are constructed using the same sequential linking approach, enabling the computation of per capita measures that preserve growth dynamics over the full historical span.

**Final coverage.** The database includes a broad set of macroeconomic variables spanning real activity, public finances, labor markets, and external accounts. In addition to real GDP, key variables used in our analysis include government revenues and expenditures, public debt, investment, unemployment, and measures of human capital. The breadth of coverage varies across variables, but in most cases the GMD substantially extends both the temporal depth and cross-country coverage available in existing datasets.

To illustrate this, Appendix Figure A.1 compares the coverage of the GMD to the next-most comprehensive alternative, showing that the database provides a larger number of country-year observations and significantly longer time series for all variables. A complete list of variables, sources, and coverage is provided in the Technical Appendix.

## 2.2 New financial crisis chronology

Our analysis examines the long-run effects of banking crises using newly a harmonized chronology constructed from the principal datasets in the literature: Reinhart and Rogoff (2009*b*); Laeven and Valencia (2013); Jordà, Schularick and Taylor (2017); Baron, Verner and Xiong (2021). These sources differ in definitions, empirical criteria, and coverage, but collectively represent the most comprehensive record of historical banking crises.

We begin by taking the union of all crisis episodes identified across these sources, ensuring that the initial set of candidate episodes is as comprehensive as possible. We then reconcile discrepancies across sources through a systematic review of historical evidence. For each country-year flagged as a crisis in at least one source but not others, we examine contemporaneous accounts of banking system distress, including documented bank failures, widespread insolvencies, government interventions, or disruptions to financial intermediation. Episodes characterized solely by depositor runs without broader impairment of the banking system are excluded.

Differences in the dating of crises across sources are resolved by defining the onset year as the earliest year recorded for a given episode. When crises span multiple years, we likewise record the first year as the onset. This approach avoids mechanically fragmenting single events while preserving consistency in the timing of crisis shocks.

The resulting indicator, $BankingCrisis$, equals one in the onset year of each crisis episode and zero otherwise. In the final sample, approximately 1.8% of country-year observations contain a banking crisis. Across countries, a new crisis occurs approximately every 35 years.

Table 1 reports summary statistics for the full sample. Panel A summarizes the main macroeconomic variables used in the analysis. The data span a broad set of countries and time periods, generating substantial variation in output levels, fiscal outcomes, and labor market conditions. This dispersion reflects the wide cross-country coverage and long historical horizon of the dataset.

Panel B reports summary statistics for crisis indicators. The mean of each indicator corresponds to the share of country-year observations in which a crisis occurs. The rows labeled "mean by country" report the average frequency with which a given country experiences the event over time, while the rows labeled "mean by year" report the cross-sectional share of countries experiencing the event in a given year. These statistics highlight that crises are infrequent in any given country-year but recur over time within countries and, at times, affect multiple countries simultaneously.

## 2.3 Additional data

**Other crisis indicators.** To distinguish banking crises from other forms of macroeconomic instability, we incorporate indicators for sovereign debt crises and currency crises from Reinhart and Rogoff (2009*b*).

Table 1: Summary Statistics

| **Panel A: Continuous Variables** | | | | | | |
|---|---|---|---|---|---|---|
| | Mean | SD | p25 | Median | p75 | N |
| Real GDP (USD bn, 2015 prices) | 176.03 | 925.05 | 2.64 | 12.46 | 66.89 | 17,911 |
| Government Debt (% of GDP) | 51.45 | 46.67 | 20.41 | 39.72 | 67.24 | 13,629 |
| Government Deficit (% of GDP) | -2.64 | 11.96 | -4.40 | -1.80 | 0.07 | 12,419 |
| Government Revenues (% of GDP) | 22.30 | 15.81 | 11.40 | 18.99 | 29.78 | 13,491 |
| Unemployment Rate (%) | 7.83 | 6.26 | 3.52 | 6.20 | 10.05 | 7,116 |
| **Panel B: Indicator Variables** | | | | | | |
| | Mean | SD | p25 | Median | p75 | Max |
| Banking Crisis | 1.82 | 13.35 | 0.00 | 0.00 | 0.00 | 100.00 |
| Mean by country | 2.06 | 2.12 | 0.68 | 2.08 | 2.71 | 25.00 |
| Mean by year | 1.63 | 3.02 | 0.00 | 0.00 | 2.45 | 21.33 |
| Sovereign Debt Crisis | 1.46 | 11.99 | 0.00 | 0.00 | 0.00 | 100.00 |
| Mean by country | 1.42 | 1.65 | 0.00 | 0.46 | 2.29 | 6.25 |
| Mean by year | 1.39 | 2.23 | 0.00 | 0.63 | 1.49 | 19.40 |
| Currency Crisis | 2.79 | 16.47 | 0.00 | 0.00 | 0.00 | 100.00 |
| Mean by country | 2.83 | 2.05 | 1.83 | 2.29 | 4.08 | 14.00 |
| Mean by year | 2.69 | 3.44 | 0.00 | 1.43 | 4.29 | 18.57 |

*Notes:* This table reports summary statistics for the full sample. Panel A presents continuous variables, including real GDP (in billions of USD at constant 2015 prices), fiscal variables (as % of GDP), and the unemployment rate (in percent). Panel B reports indicator variables in percent. Indicator variables equal one in years in which the corresponding event occurs. "Mean by country" denotes the within-country frequency of the event, while "mean by year" denotes the cross-sectional share of countries experiencing the event in a given year.

**Education and human capital.** We measure human capital using educational attainment data from Barro and Lee (1996), which provides internationally comparable measures of schooling for the population aged 15 and above at five-year intervals beginning in 1820.

**Conflict and geopolitical shocks.** We incorporate data on interstate and civil conflicts from the Correlates of War (COW) project (Sarkees and Wayman, 2010).

**Technology adoption.** To capture cross-country differences in technological diffusion, we incorporate data from the Cross-country Historical Adoption of Technology (CHAT) dataset (Comin and Hobijn, 2009).

Together, these datasets provide complementary information on macroeconomic shocks and structural characteristics that are not directly captured in the Global Macro Database.

# 3 The Very Long-Run Output Costs of Financial Crises

## 3.1 Methodology

We estimate the effects of financial crises on real economic activity using a local projections (LP) framework in the spirit of Jordà (2005). Our baseline specification traces the cumulative response of real GDP at horizon $h$ to the onset of a banking crisis while conditioning on a long history of past output dynamics.

This approach is well suited to our setting for several reasons. First, it allows us to trace the impulse response of output to crisis events over very long horizons without imposing strong parametric restrictions on the underlying data-generating process. Second, it accommodates the unbalanced panel structure, heterogeneous crisis timing, and pronounced low-frequency dynamics that characterize global historical macroeconomic data. Third, horizon-by-horizon estimation avoids dynamic misspecification concerns associated with long-horizon VARs and yields valid inference even when crisis effects unfold gradually over time (Ramey and Zubairy, 2018; Plagborg-Møller and Wolf, 2021).

For each horizon $h$, we estimate the following specification:

$$\Delta \log Y_{i,t+h} = \beta_h \, \text{BankingCrisis}_{i,t} + \alpha_i^{(h)} + \gamma^{(h)} \log Y_t^{World} + \sum_{\ell=1}^{L} \phi_\ell^{(h)} \Delta \log Y_{i,t-\ell} + \varepsilon_{i,t}^{(h)}, \qquad (1)$$

in which $i$ indexes countries and $t$ indexes years. The crisis indicator, $\text{BankingCrisis}_{i,t}$, equals one in the first year of a crisis episode, as defined in Section 2. The dependent variable $\Delta \log Y_{i,t+h} \equiv \log Y_{i,t+h} - \log Y_{i,t-1}$ measures the cumulative change in real GDP from the year preceding the crisis through horizon $h$, so that the coefficients $\beta_h$ can be interpreted as impulse responses relative to the pre-crisis trajectory. Country fixed effects $\alpha_i^{(h)}$ absorb time-invariant differences in long-run income levels. We control for global output through $\log Y_t^{World}$ to account for common global shocks. Standard errors are heteroskedasticity-robust and clustered at the country level.

Because output dynamics are highly persistent at low frequencies, and financial crises tend to occur following extended expansions rather than as purely transitory shocks, a central feature of our empirical strategy is the inclusion of long lags of output growth. These lags are denoted by $\sum_{\ell=1}^{L} \phi_\ell^{(h)} \Delta \log Y_{i,t-\ell}$, with lag lengths $L \in \{5, 10, 25\}$ across specifications. Without conditioning on sufficiently long pre-crisis histories, estimated impulse responses risk conflating the causal effect of crises with the unwinding of pre-existing trends. This concern is especially acute at multi-decade horizons, in which slow-moving structural forces such as secular growth slowdowns, institutional change, or demographic transitions could otherwise be attributed to the crisis event itself. By controlling flexibly for decades of prior output growth, the specification ensures that identifying

variation comes from deviations from deeply embedded country-specific trajectories rather than cross-country differences in long-run growth regimes.

The ability to include such long lag structures while estimating responses at multi-decade horizons relies on the scale of the GMD. The combination of broad country coverage and long historical time series provides sufficient statistical power to estimate local projections with extensive lag structures without sacrificing precision.

**Interpretation.** The estimated coefficients $\{\beta_h\}_{h=0}^{50}$ trace out the average path of real GDP following a banking crisis relative to the counterfactual of no crisis, conditional on the long pre-crisis output history. The identifying assumption is that, after conditioning on long lags of output growth, global output dynamics, and country-specific characteristics, the timing of crisis onset is orthogonal to unobserved determinants of future output at horizons $h = 1, \ldots, 50$. We show below that output dynamics in the twenty years prior to crisis onset display no systematic pre-trends once conditioning on lagged output growth, which supports the validity of this assumption. However, fully isolating exogenous variation in crisis timing remains challenging even with extensive controls. We therefore interpret the estimates as capturing the persistent output losses associated with financial crises rather than strictly causal effects.

### 3.2 Baseline

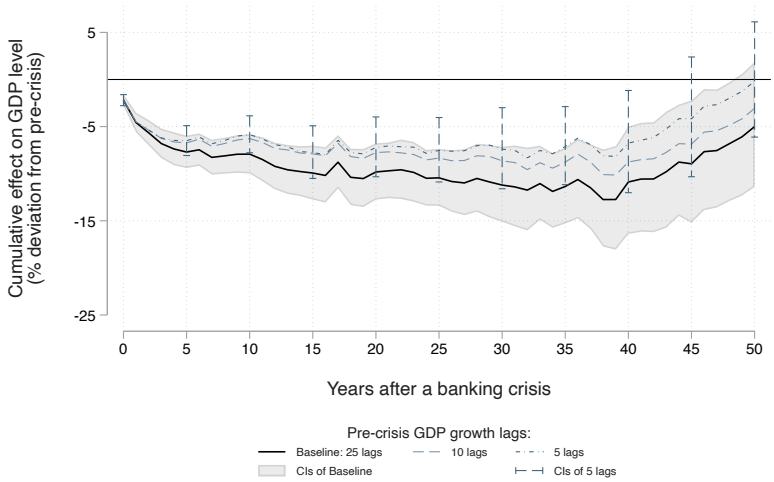
We estimate equation (1) using local projections. Our baseline uses the full panel of available data to maximize statistical power. Because the panel is unbalanced, the set of crisis observations varies across horizons as longer forward observations become unavailable for more recent crises. As a result, estimates at longer horizons place greater weight on earlier historical episodes.

To assess whether this feature affects the estimated long-run responses, we also report results from a balanced panel that restricts attention to crisis episodes with at least fifty years of post-crisis data, in addition to the years required to construct the lagged control variables. This restriction ensures that the composition of crisis episodes remains constant across horizons, allowing for a direct comparison of short- and long-run dynamics within the same sample.

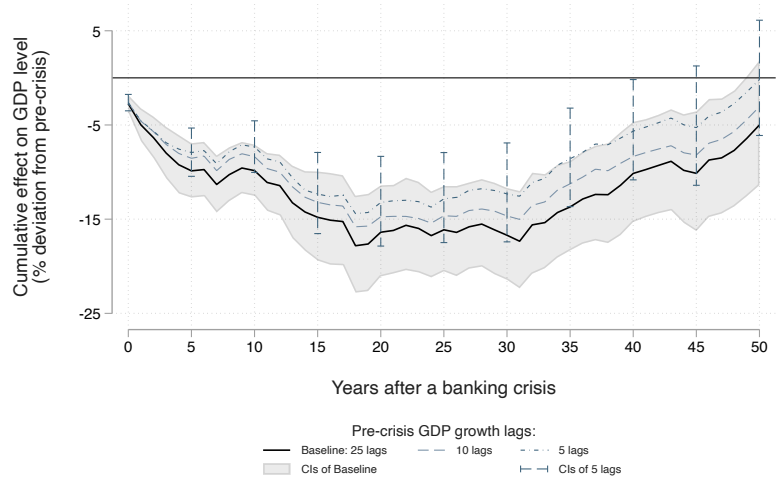
Figure 2 presents our baseline estimates of the impulse responses of cumulative real GDP following a banking crisis. The horizontal axis plots years since crisis onset, while the vertical axis reports cumulative percent deviations of real GDP relative to its pre-crisis level. The dark solid line corresponds to the baseline specification with 25 lags of pre-crisis growth, and the shaded gray region denotes the associated 90 percent confidence interval. The lighter dashed lines report alternative specifications with shorter lag lengths ($L = 10$ and $L = 5$), with confidence intervals shown selectively to preserve visual clarity.

Figure 2: Long-Run Impact of Banking Crises

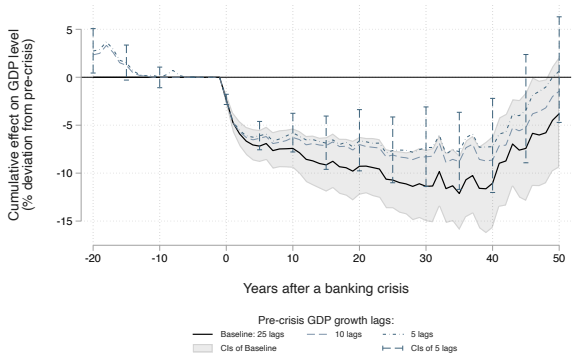
(a) Full Panel



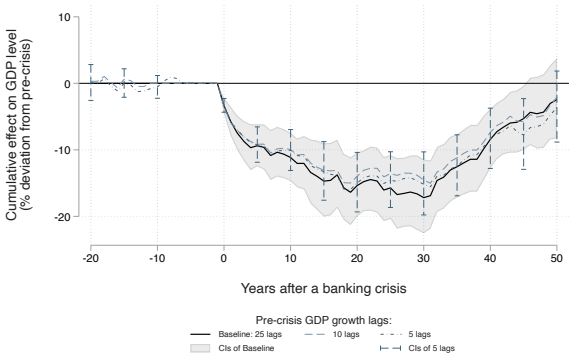

(b) Balanced Panel




(c) Pre-Trends: Full Panel




(d) Pre-Trends: Balanced Panel




*Notes:* This figure reports local projection estimates of the cumulative response of real GDP to a banking crisis, relative to the year prior to the crisis. The horizontal axis denotes years since the crisis onset ($h = 0$), and the vertical axis plots cumulative percent changes in real GDP relative to the pre-crisis level. Estimates are obtained from equation (1). Panel (a) uses the full panel of available data, allowing the set of crisis observations to vary across horizons. Panel (b) restricts attention to a balanced panel of crises for which the full set of leads and lags is available. Panels (c) and (d) extend these specifications to include horizons up to twenty years prior to the crisis. Each line corresponds to a specification that conditions on $L$ lags of pre-crisis GDP growth, with $L \in 5, 10, 25$, where available. Shaded areas denote 90 percent confidence intervals for $L = 25$ while vertical CIs are shown every 5 years for $L = 5$. Standard errors are clustered by country.

Conditioning on longer histories of pre-crisis growth modestly increases the magnitude of estimated losses at all horizons. This pattern is consistent with the fact that financial crises tend to occur following extended expansions (e.g., Reinhart and Rogoff, 2009*b*; Schularick and Taylor, 2012). With shorter lag structures, the specification does not fully account for these pre-crisis dynamics, so the counterfactual path implicitly incorporates part of the subsequent slowdown. Including longer lag histories more completely absorbs these slow-moving trends, implying a stronger pre-crisis trajectory and therefore a larger estimated deviation following the crisis.

Figure 2b reports results from the balanced panel. Estimated losses are larger in magnitude, with trough declines of approximately 15 percent, but the overall shape of the impulse response is very similar to that in the full panel. This difference is informative about sample composition. Because the balanced panel restricts attention to crisis episodes with at least fifty years of post-crisis data, it places relatively greater weight on earlier historical crises and excludes more recent episodes for which long-run outcomes are not yet observed. The larger losses in the balanced panel therefore suggest that the earlier crises that dominate this balanced sample are, on average, associated with more severe long-run scarring. At the same time, the similarity in the overall response pattern across the two samples indicates that the central result—large and highly persistent output losses following banking crises—is not driven by horizon-specific changes in sample composition.

Figures 2c and 2d provide a diagnostic for potential pre-trends by extending the local-projection specification to horizons prior to crisis onset for the full and balanced panels, respectively. For negative horizons $h < 0$, the dependent variable remains defined as $\Delta \log Y_{i,t+h} \equiv \log Y_{i,t+h} - \log Y_{i,t-1}$, so that it measures cumulative output growth from year $t + h$ up to year $t - 1$. For horizons close to the crisis date, these pre-crisis outcomes are mechanically collinear with the lagged output-growth controls already included in equation (1), so we only estimate these specifications with $L = 5$ and $L = 10$ lags.

In the full panel (2c), output rises prior to crisis onset, consistent with the view that banking crises are often preceded by sustained expansions. In the balanced panel (2d), by contrast, the estimated pre-crisis coefficients remain close to zero throughout the identified horizons. Because the balanced panel holds the set of crisis episodes fixed across horizons, variation in the estimated pre-crisis coefficients reflects within-episode dynamics rather than changes in sample composition, providing a cleaner diagnostic for pre-trends. The absence of pre-trends in this setting indicates that, conditional on the long history of lagged output growth included in the baseline specification, there is no systematic evidence of differential pre-crisis dynamics.[1]

[1]Because local projections are estimated separately at each horizon, the impulse responses for $h \geq 0$ would be identical across specifications if the estimation sample were unchanged. However, extending the specification to include negative horizons requires restricting the sample to observations with sufficiently long pre-crisis histories. As a result, the balanced-panel estimates that include pre-trends are based on a slightly different sample than those reported in Figure 2b, which leads to small differences in the estimated impulse responses.

### 3.3 Robustness

We provide several robustness checks of the baseline results. Across all exercises, the central finding of large and persistent output losses following financial crises remains unchanged.

In all figures, we present the baseline specification using the full panel and conditioning on twenty-five lags of pre-crisis output growth. We show these baseline results as a dark line, with shaded gray bands denoting the corresponding confidence intervals. Alternative specifications are reported as thinner lines without confidence intervals, and in selected cases we report confidence intervals at five-year intervals for specifications that include the full set of controls. Appendix C.2 reports the full set of results across alternative lag structures and sample restrictions.

**Controlling for contemporaneous shocks and macroeconomic conditions.** Financial crises often coincide with other major macroeconomic disruptions, raising the concern that estimated output losses attributed to banking crises may partly reflect the effects of concurrent shocks (e.g., Kaminsky and Reinhart, 1999; Laeven and Valencia, 2013). In particular, wars and other forms of macro-financial distress such as currency crises and sovereign debt crises are frequently clustered in time and may themselves have large and persistent effects on real economic activity. If such shocks are correlated with crisis timing, failing to account for them could bias the estimated dynamic response to banking crises.

To address this concern, we augment the baseline local projections specification by controlling for additional contemporaneous shocks at the crisis onset year. Specifically, we first include an indicator for wars in year $t$, and then expand the conditioning set to include indicators for currency crises and sovereign debt crises, as defined in Section 2. These controls enter contemporaneously alongside the banking crisis indicator, so that identification continues to rely on deviations from pre-crisis output trajectories rather than post-crisis dynamics.

Figure 3a reports the resulting impulse responses from controlling for contemporaneous wars and other types of crises while Figure 3b instead examines robustness to conditioning on a broader set of macroeconomic variables, including inflation, interest rates, credit aggregates (M3), investment, the current account balance, government debt, and the real effective exchange rate. In each panel, the dark line corresponds to the baseline specification, while thinner lines report estimates from specifications that add each control individually. Confidence intervals for the specification including all controls are reported at five-year intervals.

Across these specifications, the estimated output dynamics following banking crises remain highly stable. Figure 3a shows that there are imperceptible differences relative to the baseline while Figure 3b shows that conditioning on additional macroeconomic variables does not materially affect the magnitude, timing, or persistence of the estimated effects.

**Excluding major global disruptions.** While the preceding results address the role of contemporaneous shocks and macroeconomic conditions, a related concern is that long-run estimates may be driven by a small number of extreme global events. In particular, the two World Wars and the Great Depression represent tail events of economic disruption that affected a large share of countries simultaneously and are often treated separately in the historical macroeconomic literature.

To address this concern, we re-estimate the baseline specification after excluding crisis episodes associated with these major global disruptions. We conduct three exercises. First, we exclude crises occurring during World War I (1914–1919) and World War II (1939–1944). Second, we exclude crises corresponding to the Great Depression period (1929–1939). Third, we exclude all of these episodes simultaneously. In each case, the sample retains the full panel of countries, but crisis indicators during these periods are set to zero.

Figure 3c reports the corresponding impulse responses. Across all exercises, the estimated dynamics closely resemble the baseline results, indicating that the average effects we find do not depend on these global shocks. If anything, estimated losses are more persistent when these episodes are excluded, consistent with the fact that post-war and post-Depression recoveries involve strong output growth that, when included, is attributed to crisis episodes and thus attenuates the estimated persistence of losses.

**Alternative crisis chronologies.** A final concern relates to the construction of the crisis chronology itself. Our baseline crisis indicator is a harmonized chronology that aggregates information across the principal banking-crisis datasets used in the literature and resolves differences in timing and episode definitions through a transparent consolidation procedure (Section 2). While this approach is designed to maximize coverage and comparability over long horizons, it is natural to ask whether our results depend on the specific construction of the crisis chronology.

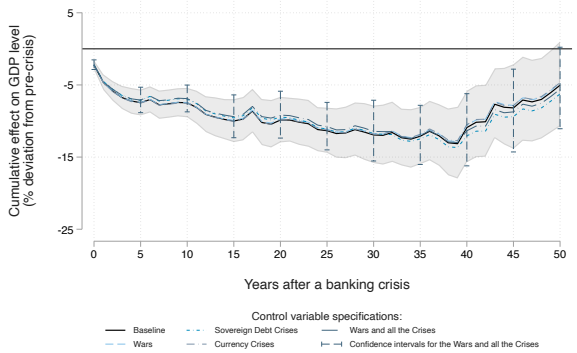
To assess this, we re-estimate equation (1) using banking-crisis indicators taken directly from each underlying source dataset, replacing $BankingCrisis_{i,t}$ with the source-specific crisis onset indicator while holding fixed the outcome definition, control set, lag structure, and sample restrictions. This exercise isolates the sensitivity of the estimated impulse responses to alternative crisis dating conventions and episode definitions.

Figure 3d shows that the estimated output response is highly robust across crisis chronologies. Regardless of whether crises are dated using Reinhart and Rogoff (2009*b*), Jordà, Schularick and Taylor (2017), Baron, Verner and Xiong (2021), or their union, banking crises are followed by large and persistent output losses that remain substantial for multiple decades. Differences across chronologies are concentrated at short horizons and reflect variation in crisis timing and thresholds across sources, but the central finding of long-run scarring is unchanged.

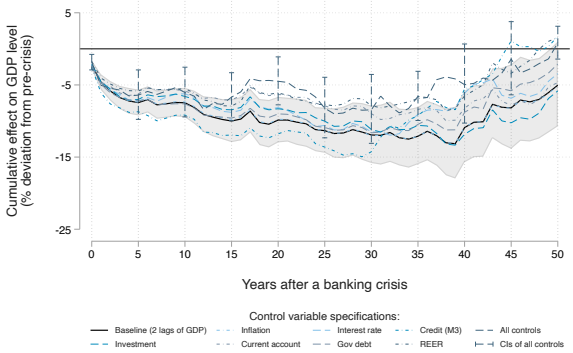
Taken together, these robustness exercises show that the estimated long-run output losses fol-

Figure 3: Robustness

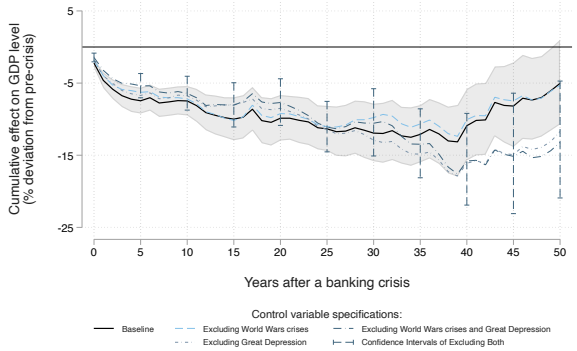
(a) Wars, Conflicts, and Other Crises



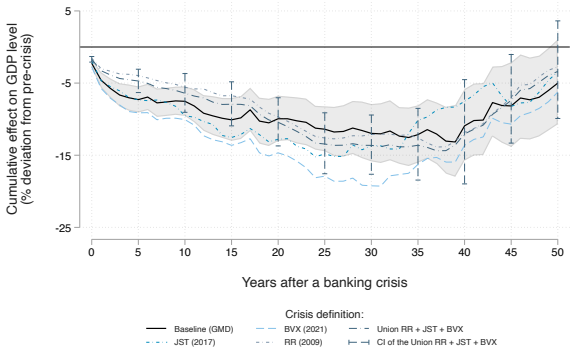

(b) Additional Macroeconomic Conditions




(c) Excluding World War & Great Depression




(d) Alternative Crisis Chronologies




*Notes:* This figure reports local projection estimates of the cumulative response of real GDP to a banking crisis, relative to the year prior to the crisis. The horizontal axis denotes years since crisis onset ($h = 0$), and the vertical axis plots cumulative percent deviations of real GDP from the pre-crisis level. Estimates are obtained from equation (1) using the full panel and conditioning on 25 lags of pre-crisis GDP growth. In all panels, the dark solid line corresponds to the baseline specification, with the shaded region denoting 90 percent confidence intervals based on standard errors clustered at the country level. Confidence intervals for specifications including the full set of controls are shown at five-year intervals. Panel (a) augments the baseline specification with controls for wars and conflicts (including interstate," intrastate," and "extrastate" conflicts from Sarkees and Wayman (2010)) as well as sovereign debt and currency crises. Panel (b) instead augments the specification with macroeconomic controls measured at the time of the crisis, including inflation, interest rates, credit (M3), investment, the current account balance, government debt, and the real effective exchange rate. Panel (c) excludes crisis episodes occurring during World War I, World War II, and the Great Depression. Panel (d) reports estimates using alternative banking crisis chronologies. In each panel, thinner lines correspond to specifications that modify the baseline by adding one control or using an alternative definition.

lowing banking crises are not driven by contemporaneous macroeconomic shocks, extreme global events, or the particular construction of the crisis chronology. Across a wide range of alternative specifications, samples, and definitions, the impulse responses remain remarkably stable.

### 3.4 Placebo Analysis and Alternative Estimators

To further assess whether our results depend on the precise specification of the baseline local projections, we conduct several additional sets of exercises. First, we implement a placebo test that randomizes crisis timing while preserving the overall frequency of crises in the data. Second, we compare our baseline estimates to several alternative estimators and outcome definitions, including

local projections in levels and long differences, and the mean group estimator.

**Placebo exercise.** A potential concern with long-horizon local projections is that persistent output dynamics may mechanically generate spurious treatment effects when shocks occur at particular points along slow-moving macroeconomic trends. While our long lag structure is designed to absorb such dynamics, we also conduct a placebo exercise based on simulated crisis timing.

Specifically, we generate placebo crisis indicators that preserve the overall frequency of crises observed in the data but randomize their occurrence across country-year observations. For each simulation, crisis onsets are drawn independently so that the unconditional probability of a crisis matches the sample mean of the banking crisis indicator. We then re-estimate the baseline local projections specification in equation (1) using each simulated crisis series. Repeating this procedure one hundred times yields the empirical distribution of impulse responses that would arise if crisis timing were unrelated to economic conditions. If the estimated effects were driven purely by underlying output persistence or other low-frequency trends, randomly assigning crises with the same overall frequency would generate similar impulse responses.

Figure 4a shows that the impulse responses generated by randomly assigned crises are centered near zero at all horizons and display no persistent decline in output. This pattern indicates that the baseline results are unlikely to arise from spurious correlations between crisis timing and slow-moving macroeconomic dynamics.

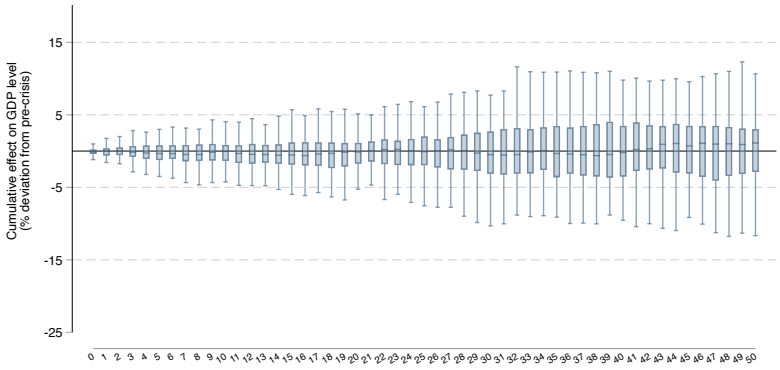
**Alternative LP specifications.** Our baseline specification measures the outcome as the cumulative change in log real GDP relative to the year prior to the crisis. To assess whether the results depend on this normalization, we estimate alternative local-projection specifications using (i) the level of log real GDP and log real GDP per capita, and (ii) long differences in log real GDP and log real GDP per capita.

These specifications provide complementary views of the same underlying dynamics. Estimating LPs in levels makes the post-crisis divergence in output paths directly visible, while long-difference specifications reduce sensitivity to short-run fluctuations and place greater weight on medium- and long-run movements. Figure 4b shows that all alternative specifications closely track the baseline impulse response with no substantive differences in dynamics. Overall, the central pattern of long-run scarring is robust to these alternative outcome definitions.

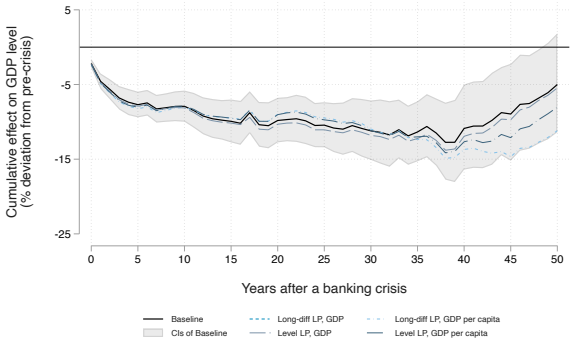
**Mean group estimator.** We next re-estimate the dynamic effects of banking crises using the mean group estimator of Pesaran and Smith (1995). Unlike pooled local projections, which impose common slope coefficients across countries, the mean group estimator allows the dynamic response to vary flexibly at the country level and then averages the resulting country-specific coefficients.

Figure 4: Placebo Exercise and Alternative Estimators

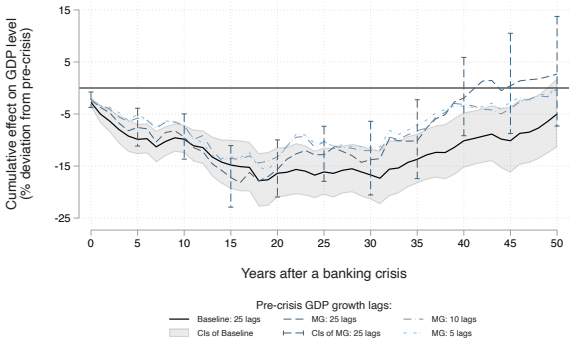
(a) Placebo Exercise




(b) Alternative LP Specifications




(c) Mean Group Estimator




*Notes:* This figure reports additional robustness exercises for the estimated long-run effects of banking crises on real GDP. The horizontal axis denotes years since the crisis onset ($h = 0$), and the vertical axis plots cumulative percent deviations of real GDP relative to the pre-crisis level. Panel (a) reports the distribution of impulse responses obtained from placebo simulations in which crisis timing is randomized while preserving the overall frequency of crises. In panels (b)–(d), the dark line denotes the baseline local-projection estimate using the full panel and conditioning on 25 lags of pre-crisis GDP growth, and the shaded region denotes the corresponding 90 percent confidence interval (standard errors clustered at the country level). Alternative specifications are shown as thinner lines without confidence intervals. Panel (b) compares the baseline estimates to alternative local-projection specifications using levels and long differences of log real GDP per capita. Panel (c) reports estimates based on the mean group estimator of Pesaran and Smith (1995) using the balanced panel.

This estimator is informative in our setting because it assigns equal weight to each country's estimated response and therefore allows for the possibility that crisis effects differ systematically across countries. We implement it using the same controls, lag structure, and outcome normalization as in the baseline analysis.

To isolate the role of cross-country heterogeneity, we implement this estimator on the balanced panel. This ensures that differences relative to the baseline estimates reflect the relaxation of the homogeneity assumption rather than changes in sample composition across horizons. In contrast, applying the mean group estimator to the unbalanced panel would introduce additional variation from the changing set of countries contributing to the estimates at longer horizons.

Figure 4c shows that the mean group estimates continue to display large and persistent output

losses following banking crises. The qualitative dynamics closely mirror those in the pooled baseline, with output remaining substantially below its pre-crisis trajectory for several decades. However, the estimated effects are somewhat smaller in magnitude than in the pooled specification.

This pattern indicates that the pooled estimates place relatively greater weight on episodes with more severe and persistent downturns. By assigning equal weight to each country, the mean group estimator reduces the influence of such episodes and yields a more moderate estimate of the average response. Taken together, these results suggest that while the magnitude of crisis-induced losses varies across countries, the persistence of output declines is a robust feature of the data. We examine the sources of this heterogeneity in Section 4.

Taken together, these exercises indicate that the persistent output losses following banking crises are not driven by spurious low-frequency trends or by the particular econometric structure of the baseline specification. The placebo results show that randomly assigned crises do not generate persistent output declines, while alternative outcome definitions and estimators yield qualitatively similar impulse responses. We therefore interpret the long-run output losses documented above as a robust feature of the data rather than an artifact of specification. We next turn to the mechanisms that may account for this persistence.

# 4 Mechanisms

The preceding analysis establishes that financial crises are followed by exceptionally large and persistent output losses. This section investigates the mechanisms underlying this persistence. Our approach is guided by the view that crises operate through disruptions to financial intermediation, which tighten balance sheet constraints and impair the allocation of capital. Rather than testing a single structural model, we examine several empirically distinct manifestations of this mechanism.

We focus on four complementary channels. First, we examine whether financial crises increase vulnerability to subsequent crises, which would reflect persistent fragility in financial systems and could compound output losses over time. Second, we study the evolution of technoloical adoption and capital accumulation, including both physical investment and human capital, to assess whether tighter financial constraints translate into sustained reductions in productive capacity. Third, we analyze the response of government balance sheets, examining how revenues, expenditures, and public debt evolve following crises. Finally, we investigate heterogeneity in these dynamics by distinguishing between global and non-global crises and between advanced and emerging economies, allowing us to assess how the severity and persistence of output losses vary with the availability of external and fiscal financing.

Taken together, these exercises provide a unified perspective on persistence: financial crises generate long-run economic scarring because they disrupt financial intermediation along multiple

margin, leading to sustained constraints on investment, adjustment, and recovery.

### 4.1 Crisis Recurrence and Financial Fragility

A natural manifestation of tighter financial constraints following a crisis is increased vulnerability to subsequent episodes of financial distress. If crises weaken financial institutions, erode confidence, or degrade regulatory and fiscal capacity, they may leave the financial system persistently fragile (e.g., Reinhart and Rogoff, 2009*b*). In this case, crises may cluster over time, compounding output losses and contributing to slow and incomplete recoveries. This channel is particularly relevant in our setting, in which output remains depressed for several decades following the initial crisis.

To assess whether financial crises increase the risk of subsequent crises, we estimate the dynamic response of the cumulative probability of experiencing an additional crisis using a local projections framework analogous to that for output. Specifically, at horizon $h$, the outcome is an indicator equal to one if a country experiences at least one additional banking crisis between years $t+1$ and $t+h$, and zero otherwise. We then project this cumulative indicator on the onset of a banking crisis at time $t$, conditioning on country fixed effects and the same long history of pre-crisis output dynamics as in the baseline specification.[2]

Although the dependent variable is cumulative by construction, the estimated coefficients need not be monotonic because they capture conditional differences in crisis likelihood rather than cumulative averages. In addition, because crises often cluster and the dependent variable becomes saturated at longer horizons, estimated effects may fluctuate rather than rise steadily.

Figure 5a reports the estimated effect of a banking crisis on the probability of experiencing at least one additional crisis within the subsequent $h$ years. In the short run, the estimated effects are small or negative, reflecting the fact that crises are dated by their initial onset and that it is less likely for a distinct new crisis episode to occur while the initial crisis is still unfolding. At longer horizons, however, the estimated effect rises sharply, indicating that countries experiencing a banking crisis face a substantially higher probability of subsequent crises over the following decades.

The effect peaks at intermediate horizons and then gradually declines, reflecting both the clustering of crises in the medium run and the saturation of the cumulative outcome at longer horizons for all countries. Throughout, the estimates remain economically meaningful and statistically significant over a wide range of horizons.

Taken together, these results provide evidence of persistent financial fragility following banking

[2] Formally, for each horizon $h$ we estimate

$$Y_{i,t,h} = \beta_h \, \text{BankingCrisis}_{i,t} + \alpha_i^{(h)} + \gamma^{(h)} \log Y_t^{\text{World}} + \sum_{\ell=1}^{L} \phi_\ell^{(h)} \Delta \log Y_{i,t-\ell} + \varepsilon_{i,t}^{(h)},$$

where $Y_{i,t,h}$ is an indicator equal to one if country $i$ experiences at least one additional banking crisis between years $t+1$ and $t+h$, and zero otherwise.

crises. While the likelihood of additional crises eventually diminishes, the elevated risk over extended horizons suggests that crisis episodes are not isolated events but part of broader periods of financial instability that can compound output losses over time. We next assess whether this channel can account for the long-run output losses documented in the baseline analysis.

#### 4.1.1 Accounting for Subsequent Crises

To evaluate the quantitative importance of crisis recurrence, we re-estimate the baseline local projections while explicitly accounting for the occurrence of crises between the initial event and each projection horizon. We implement two complementary approaches.

In the first approach, we augment the baseline specification with a control for the proximity of the most recent financial crisis prior to horizon $h$. Specifically, at each horizon we include a variable measuring the number of "years since" the most recent financial crisis as of time $t+h$. This variable captures the influence of intervening crises, allowing output dynamics to depend on the recency of financial distress rather than solely on the initial event.[3] Because this measure is undefined for countries that never experience a crisis, the specification is estimated on the subsample of countries with at least one crisis episode.

In the second approach, we directly control for crisis episodes occurring anywhere within the window $[t-L,\, t+h-1]$ by including crisis indicators for both past and future periods in the local projection specification. This absorbs the effects of any intervening crises, allowing the coefficient on the initial crisis to isolate its direct impact while conditioning on the full sequence of crisis events.[4]

Figure 5b reports the resulting impulse responses. Across both approaches, explicitly accounting for intervening crises does not attenuate the estimated long-run output losses following a banking crisis as output remains substantially below its pre-crisis trajectory for several decades, with little evidence of convergence. If anything, the estimated effects remain comparable or somewhat larger in magnitude than in the baseline specification.

Taken together, these results indicate that while financial crises increase the likelihood of subsequent crises, crisis recurrence alone cannot account for the magnitude and persistence of the long-run output losses documented in the baseline analysis. Instead, the results point to additional

[3] This specification augments equation (1) with the additional regressor $\text{YearsSince}_{i,t+h}$, which measures the number of years elapsed since the most recent financial crisis as of time $t+h$ and resets to zero at each new crisis.

[4] Leads of the crisis indicator are constructed using missing-dummy methods following Teulings and Zubanov (2014), so that the effective sample remains comparable across horizons. Formally, for each horizon $h$ we estimate

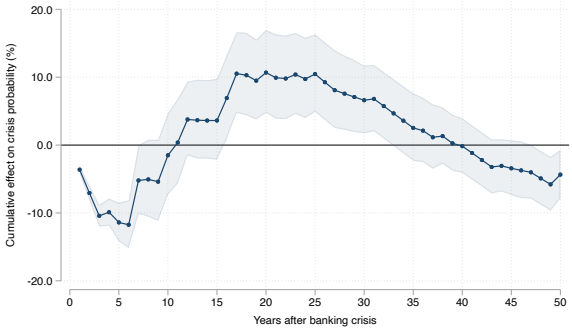
$$\Delta \log Y_{i,t+h} = \beta_h \,\text{BankingCrisis}_{i,t} + \alpha_i^{(h)} + \gamma^{(h)} \log Y_t^{\text{World}} + \sum_{\ell=1}^{L} \phi_\ell^{(h)} \Delta \log Y_{i,t-\ell} + \sum_{s=-L}^{h-1} \kappa_s^{(h)} \,\text{BankingCrisis}_{i,t+s} + \varepsilon_{i,t}^{(h)},$$

where the additional terms include $L$ lags and $h-1$ leads of the crisis indicator, thereby controlling for crisis episodes occurring anywhere within the window $[t-L,\, t+h-1]$.

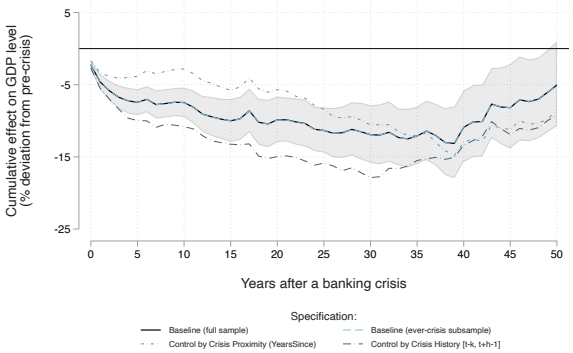
channels through which financial crises generate sustained economic scarring.

Figure 5: Subsequent Financial Crises

(a) Impact on Subsequent Financial Crises




(b) Accounting for Subsequent Crises




*Notes:* Panel (a) reports local projection estimates obtained from equation (1) of the effect of a banking crisis on the cumulative probability of experiencing a subsequent financial crisis. The dependent variable at horizon $h$ is an indicator equal to one if a country experiences at least one additional banking crisis between years $t+1$ and $t+h$, and zero otherwise. The horizontal axis denotes years since the initial crisis onset ($h=0$), and the vertical axis plots the estimated change in the cumulative probability of a future crisis relative to countries that do not experience a crisis at time $t$. Panel (b) reports specifications that account for intervening crisis episodes using two alternative approaches: controlling for the time since the most recent crisis and directly conditioning on crisis indicators within the window $[t-L,\ t+h-1]$. All specifications include country fixed effects and the same long history of pre-crisis output growth as in the baseline output regressions. The shaded region denotes 90 percent confidence intervals for the baseline specification, based on standard errors clustered at the country level. Thinner lines correspond to alternative specifications.

## 4.2 Capital Accumulation and Labor Markets

A central implication of tighter financial constraints following a crisis is a disruption to real economic adjustment along multiple margins. When banking systems are impaired, credit intermediation becomes less effective, limiting the ability of firms and households to finance investment and sustain employment (e.g., Bernanke and Gertler, 1989; Kiyotaki and Moore, 1997). As a result, financial crises may not only reduce short-run output but also depress the accumulation of physical and human capital and weaken labor markets, generating persistent effects on productive capacity.[5]

### 4.2.1 Investment & Technology

We begin by examining the evolution of investment and technology adoption following banking crises. Both are particularly sensitive to financial conditions because many capital expenditures require external financing. When banks experience balance sheet stress, the reduced lending capacity and heightened uncertainty that accompany crises may lead firms to delay, scale back, or cancel planned investment projects.

[5]We focus on the balanced panel for thse investment outcomes because long-horizon cumulative responses are particularly sensitive to horizon-specific sample composition in the unbalanced panel. Holding the set of crisis episodes fixed across horizons yields more stable and interpretable estimates.

Figure 6a plots the response of investment as a share of GDP following banking crises using the same local projection framework employed in the baseline analysis. The results show that the investment-to-GDP ratio declines sharply in the years immediately following the crisis and remains below its pre-crisis level for an extended period. By the horizon of four decades after the crisis, there appears to be recover. However, because this measure is a ratio, its dynamics reflect movements in both investment and output. In particular, the immediate decline in investment as a share of GDP indicates that investment falls by more than output in the aftermath of crises. At longer horizons, the relatively positive effects reflect the continued depression of output rather than excess investment.

Figure 6b shows that aggregate technology adoption, measured by the index from Comin and Hobijn (2009), also declines following banking crises and actually shows no recovery. These patterns indicate that financial crises disrupt the process of capital formation for many years after the initial shock. Because investment determines the evolution of the capital stock, sustained declines in investment translate into a lower path of productive capacity, providing a central channel through which financial crises generate persistent output losses.

#### 4.2.2 Labor Markets and Human Capital

Financial crises may also affect output persistence through their impact on labor markets and human capital accumulation. When firms face tighter financing constraints, they may reduce hiring, cut hours, or lay off workers in response to declining demand and limited access to credit (e.g., Mian, Sufi and Verner, 2017). Disruptions to firm balance sheets may also lead to the exit of financially constrained firms, further reducing employment.

Figure 6c presents the response of unemployment following banking crises using the same local projection framework as in the baseline analysis. We find that unemployment rises significantly in the aftermath of crises and recovers only gradually over time.

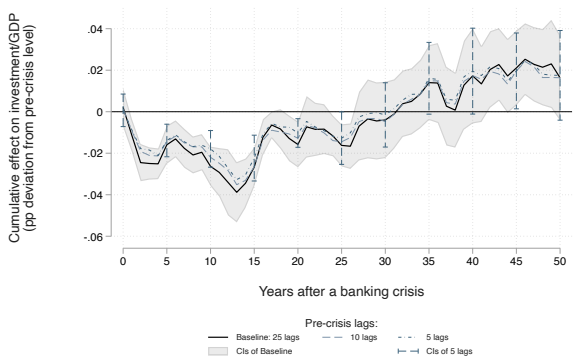
In addition to affecting employment, financial crises may disrupt investments in human capital. Households facing income shocks or tighter credit conditions may reduce spending on education or delay entry into skill-intensive occupations. At the same time, governments experiencing fiscal stress may cut education spending or reduce support for training programs.

Figure 6d presents the corresponding impulse responses for human capital accumulation (Barro and Lee, 1996). Similar to the patterns observed for employment and investment, human capital accumulation declines following crises and recovers only gradually over time. Additional proxies for human capital accumulation, including years of schooling and secondary schooling attainment, are reported in Appendix C.2 and yield similar patterns.

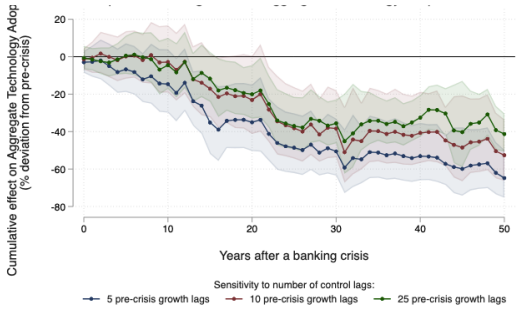
Taken together, these findings suggest that financial crises weaken both the usage and accu-

Figure 6: Capital Accumulation, Labor Markets, and Human Capital after Banking Crises

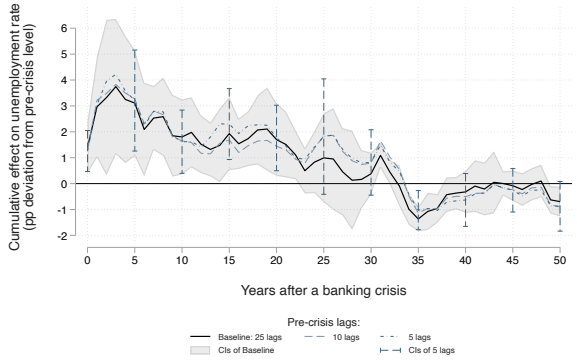
(a) Investment/GDP



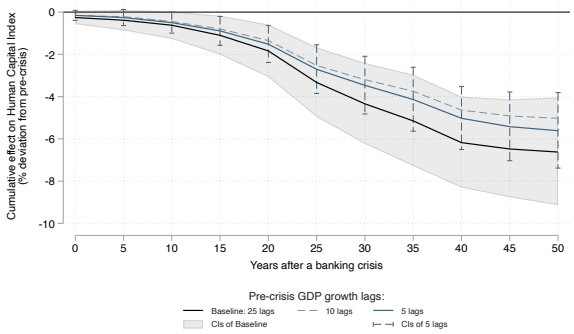

(b) Aggregate Technology Adoption




(c) Unemployment




(d) Human Capital Index




*Notes:* This figure reports local projection estimates of the response of capital accumulation, labor market outcomes, and human capital to a banking crisis. The horizontal axis denotes years since crisis onset ($h = 0$), and the vertical axis plots the cumulative response of each outcome relative to the year prior to the crisis. Estimates are obtained from equation (1) using the balanced panel and conditioning on 25 lags of pre-crisis output growth. Panel (a) reports the response of investment as a share of GDP. Panel (b) reports the response of an aggregate technology adoption index constructed from the CHAT database. Panel (c) reports the response of the unemployment rate, and Panel (d) reports the response of the human capital index. In all panels, the dark solid line denotes the baseline specification, and the shaded region denotes 90 percent confidence intervals based on standard errors clustered at the country level. Thinner lines correspond to alternative specifications.

mulation of labor inputs as persistent declines in employment reduce current output, while slower human capital accumulation lowers future productivity. These channels complement the evidence on physical investment by showing that financial constraints affect multiple components of productive capacity, contributing to sustained output losses even in the absence of subsequent crises.

## 4.3 Fiscal Impact

A final channel that we examine through which financial crises affect the economy operates through government balance sheets. As economic activity contracts and financial conditions deteriorate, fiscal aggregates respond endogenously through both revenue and expenditure channels.

Figure 7 presents impulse responses for government revenues, government spending, fiscal bal-

ances, and public debt, all as a share of GDP, using the same empirical framework as in the baseline.

The estimated responses show substantial and persistent deterioration in government balance sheets following crises. Government revenues decline sharply as economic activity contracts, even relative to GDP, indicating a contraction in the effective tax base beyond the mechanical effects of falling output. At the same time, expenditures increase in the years immediately following the crisis. As a result, fiscal deficits widen significantly and public debt rises persistently over subsequent decades.

These patterns are consistent with existing evidence showing that banking crises are associated with sharp deteriorations in fiscal positions and sustained increases in public debt (e.g., Reinhart and Rogoff, 2009*a*,*b*; Laeven and Valencia, 2013). The joint behavior of revenues, expenditures, and debt reflects the combined effects of reduced tax bases, automatic stabilizers, and crisis-related spending pressures, leading to a persistent increase in public indebtedness.

### 4.4 Heterogeneity in Constraints

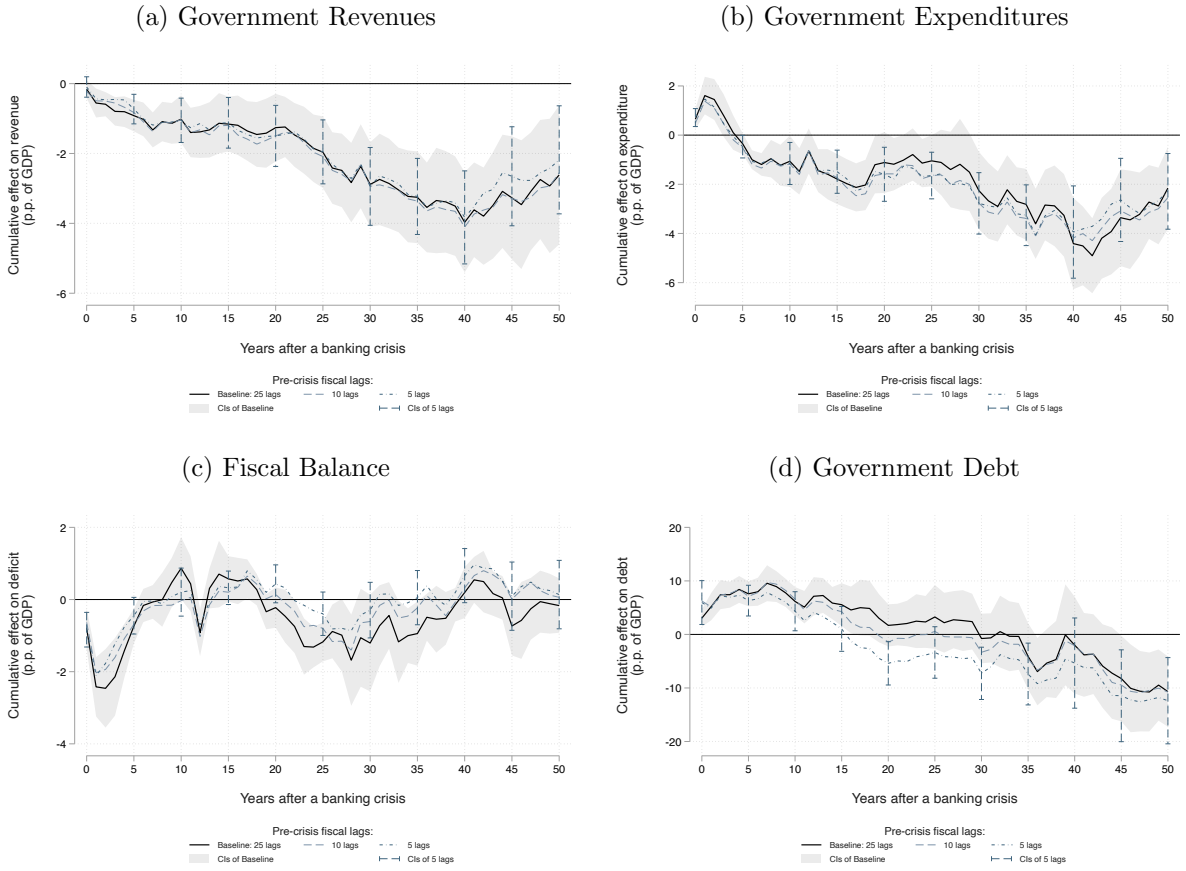
The preceding analysis documents how financial crises affect capital accumulation, labor markets, and government balance sheets. A central implication of the financial constraints framework is that these effects should vary systematically with the availability of financing. When access to external or public sources of funding is limited, the adverse effects of crises may be amplified, leading to more persistent output losses.

In the final part of the paper, we examine two dimensions along which such constraints may differ. First, we distinguish between crises that occur during periods of global financial distress and those that are more localized, capturing variation in the availability of international financing and risk-sharing. Second, we study differences across countries with varying degrees of fiscal capacity, reflecting the ability of governments to respond to shocks through their balance sheets.

#### 4.4.1 Global Banking Crises

A third dimension of financial constraints operates through the international financial system. This channel complements the evidence on capital accumulation by highlighting the role of external financing conditions in shaping recovery dynamics. When financial crises are confined to a small set of countries, capital from unaffected regions may flow toward distressed economies, facilitating bank recapitalization and supporting investment. By contrast, when crises are synchronized across many countries, or when they affect key financial centers, the global financial system itself may become impaired, limiting the availability of external financing and slowing the recovery process (Rey, 2015). In that case, disruptions to international financial intermediation may amplify and

Figure 7: Fiscal Dynamics Following Banking Crises




*Notes:* This figure reports local projection estimates of the response of fiscal outcomes to a banking crisis. The horizontal axis denotes years since crisis onset ($h = 0$), and the vertical axis plots cumulative changes in each fiscal variable relative to the year prior to the crisis, expressed as a share of GDP. Estimates are obtained from equation (1) using the full panel and conditioning on 25 lags of pre-crisis GDP growth. The dark solid line denotes the baseline specification, and the shaded region denotes 90 percent confidence intervals based on standard errors clustered at the country level. Thinner lines correspond to alternative specifications.

prolong the effects of domestic financial distress.

To investigate this mechanism, we distinguish between global and non-global banking crises using a classification that captures both the extent of international synchronization and the role of financial centers. Specifically, we compute the share of world GDP accounted for by countries experiencing a banking crisis in each year and classify crises as global when this share exceeds the 75th percentile of the historical distribution. In addition, we identify episodes in which the global financial center is itself in crisis, using the United Kingdom prior to 1945 and the United States thereafter, reflecting the transition in international financial leadership over the twentieth century (e.g., Coppola, Krishnamurthy and Xu, 2023).

Figure 8a illustrates these patterns, plotting both the share of world GDP in crisis in each year starting in 1800, and the timing of crises in the dominant financial center. The horizontal dashed lines mark the 25th percentile, median, and 75th percentile values of the share of global GDP

experiencing a crisis in any given year. We define the crises above the 75th percentile as global. We mark the crises involving the financial center are marked in blue.

Figures 8b and 8c report impulse responses separately for global and non-global crises under these two classifications. The results are broadly similar across these two definitions. They show that crises classified as global are associated with substantially more persistent output losses. By contrast, non-global crises exhibit more moderate and shorter-lived effects.

These findings are consistent with the view that international financial linkages play a key role in shaping recovery dynamics. When crises are global, the breakdown of cross-border capital flows and international financial intermediation limits the ability of countries to draw on external resources, amplifying the contraction and slowing adjustment. In contrast, when crises are more localized, access to external financing may help mitigate the severity and persistence of output losses.

#### 4.4.2 Heterogeneity in Constraints: Advanced vs Emerging Economies

A second dimension of heterogeneity in financial constraints arises from differences in countries' access to financing and ability to adjust following crises. Advanced and emerging economies differ systematically along these margins, including in the depth of financial markets, access to external borrowing, and fiscal capacity. While this classification is coarse, it provides a transparent grouping that captures differences in financial and fiscal capacity across countries.[6]

Figure 9 reports the response of real GDP to banking crises separately for advanced and emerging economies. The results show that output losses are substantially larger and more persistent in emerging economies, with the average response remaining far below that of advanced economies throughout the horizon.

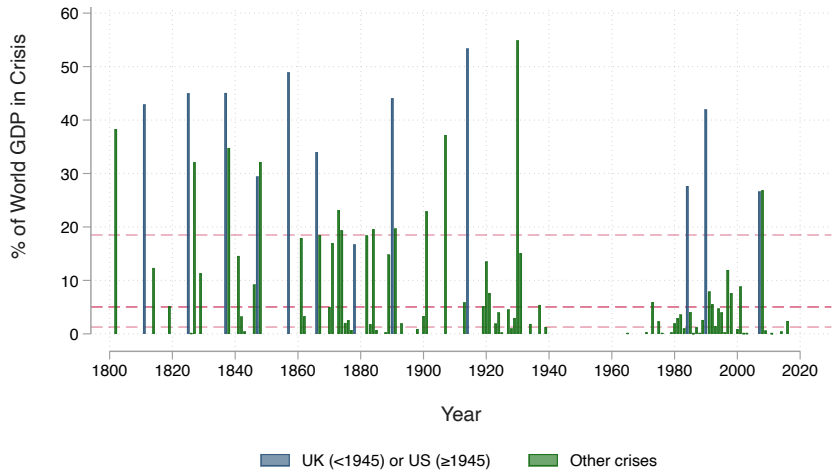
At the same time, the estimates for emerging economies exhibit considerably greater dispersion, as reflected in the wider confidence intervals. This pattern indicates that crisis outcomes in emerging markets are not only more severe on average, but also more heterogeneous across episodes.

These differences are consistent with tighter and more uneven financial constraints in emerging economies, which may limit the ability of firms, households, and governments to smooth shocks and support recovery. In particular, more limited fiscal space and weaker access to financing may restrict the scope for balance sheet adjustment following crises, contributing both to more persistent output losses and to greater variability in recovery paths.

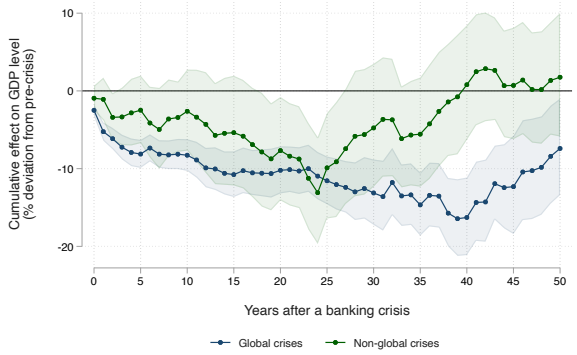
[6] Advanced economies are defined as countries belonging to the Jordà, Schularick and Taylor (2013) sample prior to 1945 and as OECD member countries thereafter. The former category includes Australia, Canada, Denmark, Finland, France, Germany, Italy, Japan, the Netherlands, Norway, Spain, Sweden, Switzerland, the United Kingdom, and the United States.

Figure 8: Global vs Non-Global Crises

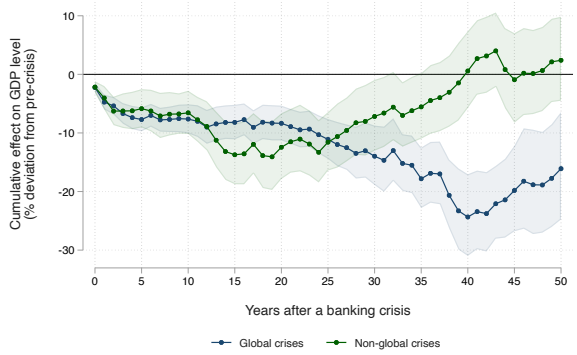
(a) Global Crises by Share of World GDP and Financial Center




(b) Global Crises by GDP Share




(c) Global Crises by Financial Center

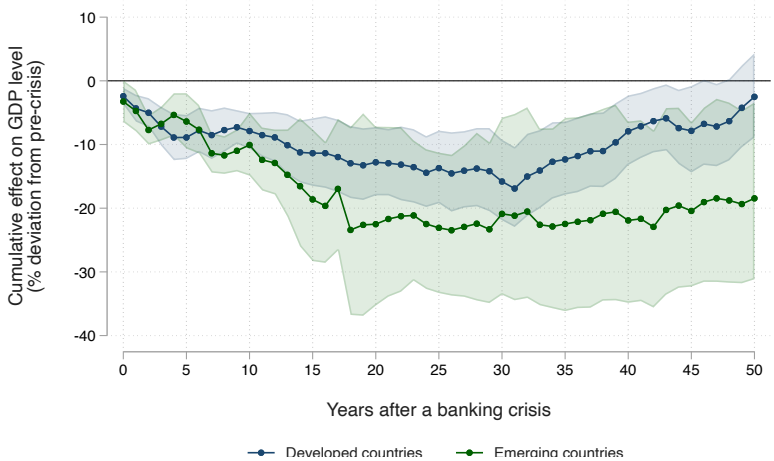



*Notes:* Panel (a) plots the share of world GDP accounted for by countries experiencing a banking crisis in a given year; crises are classified as global when this share exceeds the 75th percentile of the historical distribution. Panel (b) reports local projection estimates of the cumulative response of real GDP to banking crises under this classification, distinguishing between global and non-global crises. Panel (c) reports analogous estimates based on a financial-center classification, defining crises as global when they involve the dominant international financial center: United Kingdom prior to 1945 and United States thereafter. The horizontal axis denotes years since crisis onset ($h = 0$), and the vertical axis plots cumulative percent deviations of real GDP from the pre-crisis level. Estimates are obtained from equation (1) using the full panel and conditioning on 25 lags of pre-crisis GDP growth. In Panels (b) and (c), blue lines denote global crises and green lines denote non-global crises; shaded areas represent 90 percent confidence intervals based on standard errors clustered at the country level.

# 5 Conclusion

This paper documents that banking crises are followed by large and highly persistent output losses. Across a broad set of countries and several centuries of data, real GDP remains substantially below its pre-crisis trajectory for decades, with little evidence of full recovery. These patterns are robust across specifications, samples, and alternative crisis definitions, and are not explained by the occurrence of subsequent crises alone. Together, the results point to long-run economic scarring as a central feature of financial crises.

Figure 9: Heterogeneity in Recovery: Advanced vs Emerging Economies




*Notes:* This figure reports local projection estimates of the cumulative response of real GDP to a banking crisis, separately for advanced and emerging economies. The horizontal axis denotes years since crisis onset ($h = 0$), and the vertical axis plots cumulative percent deviations of real GDP from the pre-crisis level. Estimates are obtained from equation (1) using the full panel and conditioning on 25 lags of pre-crisis GDP growth. Blue lines denote advanced economies and green lines denote emerging economies; shaded areas represent 90 percent confidence intervals from standard errors clustered at the country level.

The evidence in this paper points to a broader interpretation of financial crises as episodes that shift the economy onto a persistently lower trajectory by tightening financial constraints. The sustained declines in investment, employment, and human capital suggest that financial frictions are not merely cyclical propagation mechanisms but a central determinant of long-run economic dynamics. In this sense, financial crises do not simply generate temporary deviations from trend; they alter the path of productive capacity itself.

Our results also highlight the importance of the international financial system as a margin of adjustment. When crises are localized, access to external financing can partially offset domestic disruptions, facilitating recapitalization and supporting investment. By contrast, when crises are globally synchronized, this channel breaks down, and the resulting contraction is both deeper and more persistent. These findings underscore that international financial integration plays a dual role: it provides a mechanism for risk-sharing in normal times, but its impairment during global crises can amplify and prolong economic downturns. Understanding the conditions under which cross-border financial flows stabilize or destabilize recovery is therefore central to assessing the macroeconomic consequences of financial integration.

Finally, the analysis emphasizes the role of fiscal capacity in shaping post-crisis adjustment. Governments can, in principle, substitute for impaired private intermediation by expanding deficits, supporting financial institutions, and sustaining demand. However, our findings indicate that this margin is highly uneven across countries. Advanced conomies with greater debt capacity appear to have the capacity for public balance sheets to some crisis-related losses, thereby partially mitigating

the contraction. In contrast, where fiscal capacity is limited, governments are less able to offset tightening financial constraints, leading to larger losses in real economic activity. This heterogeneity highlights that the effectiveness of macroeconomic stabilization is fundamentally tied to the strength of sovereign balance sheets.

# References


**Ball, Laurence.** 2014. "Long-Term Damage from the Great Recession in OECD Countries." *European Journal of Economics and Economic Policies: Intervention*, 11(2): 149–160.

**Baron, Matthew, Emil Verner, and Wei Xiong.** 2021. "Banking crises without panics." *The Quarterly Journal of Economics*, 136(1): 51–113.

**Barro, Robert J.** 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–866.

**Barro, Robert J, and Jong Wha Lee.** 1996. "International measures of schooling years and schooling quality." *The American Economic Review*, 86(2): 218–223.

**Barro, Robert J., and José F. Ursúa.** 2008. "Macroeconomic Crises since 1870." *Brookings Papers on Economic Activity*, 2008(1): 255–350.

**Bernanke, Ben S., and Mark Gertler.** 1989. "Agency Costs, Net Worth, and Business Fluctuations." *American Economic Review*, 79(1): 14–31.

**Blanchard, Olivier, Eugenio Cerutti, and Lawrence Summers.** 2015. "Inflation and Activity: Two Explorations and Their Monetary Policy Implications." *IMF Working Paper*, , (15/230).

**Cerra, Valerie, and Sweta Chaman Saxena.** 2008. "Growth dynamics: the myth of economic recovery." *American Economic Review*, 98(1): 439–457.

**Cerra, Valerie, Antonio Fatás, and Sweta Chaman Saxena.** 2023. "Hysteresis and Business Cycles." *Journal of Economic Literature*, 61(1): 181–225.

**Comin, Diego A, and Bart Hobijn.** 2009. "The CHAT dataset." National Bureau of Economic Research.

**Coppola, Antonio, Arvind Krishnamurthy, and Chenzi Xu.** 2023. "Liquidity, debt denomination, and currency dominance." National Bureau of Economic Research.

**Davis, Donald R, and David E Weinstein.** 2002. "Bones, bombs, and break points: the geography of economic activity." *American economic review*, 92(5): 1269–1289.

**Funke, Manuel, Moritz Schularick, and Christoph Trebesch.** 2016. "Going to Extremes: Politics after Financial Crises, 1870–2014." *European Economic Review*, 88: 227–260.

**Guiso, Luigi, Helios Herrera, Massimo Morelli, and Tommaso Sonno.** 2017. "Demand and Supply of Populism." *CEPR Discussion Paper*, , (11871).

**Jordà, Òscar.** 2005. "Estimation and inference of impulse responses by local projections." *American economic review*, 95(1): 161–182.

**Jordà, Oscar, Moritz Schularick, and Alan M. Taylor.** 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking*, 45(s2): 3–28.

**Jordà, Òscar, Moritz Schularick, and Alan M Taylor.** 2017. "Macrofinancial history and the new business cycle facts." *NBER macroeconomics annual*, 31(1): 213–263.

**Kaminsky, Graciela L., and Carmen M. Reinhart.** 1999. "The Twin Crises: The Causes of Banking and Balance-of-Payments Problems." *American Economic Review*, 89(3): 473–500.

**Kiyotaki, Nobuhiro, and John Moore.** 1997. "Credit Cycles." *Journal of Political Economy*, 105(2): 211–248.

**Krishnamurthy, Arvind, and Tyler Muir.** 2025. "How Credit Cycles Across a Financial Crisis." *Journal of Finance*, 80(3): 1339–1378.

**Laeven, Luc, and Fabian Valencia.** 2013. "Systemic banking crises database." *IMF Economic Review*, 61(2): 225–270.

**Mian, Atif, Amir Sufi, and Emil Verner.** 2017. "Household Debt and Business Cycles Worldwide." *Quarterly Journal of Economics*, 132(4): 1755–1817.

**Müller, Karsten, and Emil Verner.** 2024. "Credit Allocation and Macroeconomic Fluctuations." *Review of Economic Studies*, 91(6): 3645–3676.

**Pesaran, M Hashem, and Ron Smith.** 1995. "Estimating long-run relationships from dynamic heterogeneous panels." *Journal of econometrics*, 68(1): 79–113.

**Plagborg-Møller, Mikkel, and Christian K. Wolf.** 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica*, 89(2): 955–980.

**Quincy, Sarah, and Chenzi Xu.** 2026. "Regulation in Crisis: Bank Failures and Entry Barriers."

**Ramey, Valerie A., and Sarah Zubairy.** 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data." *Journal of Political Economy*, 126(2): 850–901.

**Reinhart, Carmen M, and Kenneth S Rogoff.** 2009*a*. "The aftermath of financial crises." *American Economic Review*, 99(2): 466–472.

**Reinhart, Carmen M, and Kenneth S Rogoff.** 2009*b*. *This time is different: Eight centuries of financial folly*. princeton university press.

**Rey, Hélène.** 2015. "Dilemma not trilemma: the global financial cycle and monetary policy independence." National Bureau of Economic Research.

**Romer, Christina D., and David H. Romer.** 2019. "Fiscal Space and the Aftermath of Financial Crises: How It Matters and Why." *Brookings Papers on Economic Activity*, 2019(1): 239–313.

**Sarkees, Meredith Reid, and Frank Wayman.** 2010. "COW War Data, 1816–2007 (Version 4.0)." *Correlates of War Project Data Set*, Available at https://correlatesofwar.org/data-sets/cow-war/.

**Schularick, Moritz, and Alan M. Taylor.** 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008." *American Economic Review*, 102(2): 1029–1061.

**Teulings, Coen N, and Nikolay Zubanov.** 2014. "Is economic recovery a myth? Robust estimation of impulse responses." *Journal of Applied Econometrics*, 29(3): 497–514.

**Xu, Chenzi.** 2022. "Reshaping global trade: the immediate and long-run effects of bank failures." *The Quarterly Journal of Economics*, 137(4): 2107–2161.