The Macroeconomic Consequences of Political Instability

Carlos Cesar Chavez*

June 14, 2026

Abstract

Do all coups d'état carry the same economic cost? Using a first-event-only local projection on the Bjørnskov–Rode regime panel of 180 countries from 1960 to 2022, we estimate the response of total factor productivity to coups classified by the institutional transition each one produces. The impact-year productivity contraction is 4.2 percent in coups that displace a sitting democracy and statistically negligible in failed coups and coups that leave the regime unchanged, with the pairwise contrast rejecting equality at the 5 percent level. The estimate survives country jackknife, leave-one-region-out, alternative regime-classification timing, and a wild cluster bootstrap. An option-value adoption framework organizes the asymmetry, and a two-block structural estimation matches its qualitative predictions. The economic cost of political instability concentrates in the coups that dismantle democratic institutions, so the productivity stakes of instability lie in protecting democratic survival, not in coup prevention as such.

Keywords: Coups d'état, Institutional change, Local projections, Total factor productivity, Option-value framework, Regime transitions
JEL Codes: D74, E32, O43, O47, P16

*Research Professional, Center for the Economics of Human Development, University of Chicago. Email: cchavezpadilla@uchicago.edu. All errors are my own.

1 Introduction

A coup d'état is, for the people who live through it, a sudden contraction of the future. The average effect of a coup, measured across the more than five hundred coup-year cells in the Bjørnskov–Rode regime panel, obscures an asymmetry that the recent literature has begun to recognize (Bennett et al., 2025): events differ in whether they succeed and in what they replace. A coup that displaces a sitting democracy and installs an autocratic regime is a different economic event from a coup that fails and leaves the institutional environment intact, or one that replaces one autocrat with another. The macroeconomic literature has typically estimated the average across these classes. The resulting estimate combines distinct mechanisms operating on different time horizons. This paper estimates the response by class and identifies the structural mechanism behind the asymmetry.

Using a first-event-only local projection on first-coup events in the Bjørnskov–Rode regime panel, we document that the productivity consequences of coups concentrate in transitions that replace democratic institutions with autocratic ones. The impact-year contraction in total factor productivity for democracy-destroying coups is 4.2 percent, approximately twice the average impact-year effect over all first coups in democracies. The same coefficient in coups that preserve democratic institutions, whether failed attempts or successful coups that do not change the regime type, is statistically negligible, and the pairwise contrast between the two classes survives the full robustness battery reported in Section 6.5. The joint pre-event nullity test is rejected at the 5 percent level. As developed in Appendix F.1, the pre-event gradient points in a direction that biases the impact-year contrast toward attenuation rather than inflation, and a Rambachan–Roth-style bound on the implied pre-trend leaves the contrast directionally intact.

We interpret the pattern through the option-value framework of Dixit and Pindyck (1994) developed in Section 3. Coups that displace a sitting democracy and install an autocratic regime generate large dispersion in the post-coup institutional outcome distribution, depressing irreversible investment through the option premium of waiting. Failed coups and democracy-preserving successful coups leave the post-coup institutional outcome distribution undisturbed, even when they generate substantial short-run political uncertainty about the immediate political environment. The empirical contrast between the two cells tests whether the dispersion of post-coup institutional outcomes, not the mean shift in the realized environment, is the operative pathway through which coups affect productivity. Event-level dispersion within the focal cell is itself a prediction of the framework, not a caveat against the average. Chile 1973, where the deepest impact-year TFP contraction coincides with an investment collapse, sits against Greece 1967, where a milder TFP loss pairs with an investment surge. We develop the breakdown in Section 6.4.

The structural section interprets the reduced-form pattern through an option-value adoption framework that organizes the heterogeneity into three sign-and-persistence predictions a generic "democracy-destroying coups are worse" story does not produce. First, total factor productivity contracts persistently while output growth recovers, because adoption is a stock and growth is a flow. Second, the coalition channel redirects capital toward favoured sectors, producing a positive cumulative investment response despite the negative impact-year effect. Third, the coalition channel's sign differs across the productivity and investment blocks. A two-block joint minimum-distance estimation across TFP, output growth, and gross investment lets us read each of these against the data. The Under a country-cluster bootstrap that jointly resamples the LP coefficients and the calibrated moments, the investment-block coalition parameter α_0^{inv} carries the predicted negative sign and is negative across its 90% bootstrap interval (the inference standard we adopt, its 95% interval includes zero), though imprecisely estimated and only weakly identified. The negative coalition parameter acting on a negative class-average institutional shift M_{DA} is what rationalises the positive cumulative investment response, and that rationalisation does not depend on the identification of the uncertainty-channel coefficient γ_0^{inv}. The uncertainty-channel coefficient γ_0 carries the predicted negative sign in both blocks across all empirically plausible calibrations of the dispersion moment but is not separately distinguishable from zero under fully-propagated two-step inference. Its magnitude is bounded across the J-profile range but is not sharply identified at this sample size. The investment block we report as suggestive given that the available investment aggregate combines private, government, and residential components and cannot disentangle private from state-directed investment.

The structural specification produces one further prediction we test directly: the cumulative TFP coefficient grows under the stock relation the mechanism implies, while the cumulative growth coefficient stays bounded under flow recovery. The point estimates are directionally consistent, but under the wild cluster bootstrap appropriate to the few-episode setting none of the cumulative five-year responses is distinguishable from zero (the country pairs bootstrap over-rejects with so few contributing clusters), so we read the cumulative pattern as directional and rest the inferential weight on the impact-year contrast. The point estimates, signs, and amplification ratios are robust across standard-error specifications.

The paper makes two contributions. First, it documents that aggregate averages over coup events obscure substantial heterogeneity in the events themselves, with productivity effects concentrated in democracy-destroying transitions and a clean failed-coup counterfactual that distinguishes institutional outcome dispersion from generic political uncertainty. Second, the option-value adoption framework organises this heterogeneity into the three sign-and-persistence predictions set out above, which a generic account does not produce and which the data confirm in the first two cases and find consistent in the third. Together they identify which coups carry productivity costs and recover the structural signature of the mechanism behind the asymmetry.

The remainder of the paper proceeds as follows. Section 2 situates the paper in the literature on the economic consequences of coups d'état. Section 3 develops the theoretical framework that motivates the empirical classification: a regime-uncertainty mechanism that generates distinct predictions for productivity and growth as functions of the institutional distance the coup traverses. Section 4 describes the data, the Bjørnskov–Rode coup typology, and the implications of the first-event-only restriction for sample composition. Section 5 specifies the local-projection design and the BBG-style transition decomposition. Section 6 reports the reduced-form results in three stages of increasing specificity: aggregate dynamic responses, the transition decomposition, and the focal cell with its robustness architecture. Section 7 reports the structural decomposition: single-outcome estimation, two-block joint identification, and the cumulative test. Section 8 discusses the implications.

2 Related Literature

The literature on the economic consequences of coups d'état has evolved through three identifiable phases, and the contribution of this paper is located in the third.

The first phase of the literature treated coups and the broader category of *political instability* as a single class of shock and estimated its effects on aggregate growth. Alesina and Tabellini (1990) and Alesina et al. (1996) established the canonical finding that countries with higher political instability grow more slowly, using cross-country panels and government-turnover-based measures of instability. Aisen and Veiga (2013) extended the analysis to identify mechanisms through productivity and capital accumulation. Jong-A-Pin (2009) contributed an important methodological observation: "political instability" as measured in this literature is multidimensional, and different operationalizations capture distinct phenomena (violence-driven, mass-protest-driven, within-regime-competition-driven, and regime-change-driven instability), with potentially different economic consequences. The first-phase literature, however, typically pooled across these dimensions, treating a coup, a mass protest, a government turnover, and a regime-change event as conceptually equivalent shocks.

The second phase, beginning in the 2010s, brought identification concerns to the foreground. Bjørnskov (2022) provided the pivotal result: economic conditions themselves predict coup occurrence, and the cross-country correlation between past growth and current coup probability is strong enough to confound naive estimates of the reverse effect. Bjørnskov and Rode (2020) simultaneously assembled the Regime Dataset that this paper uses, providing carefully coded coup events and pre- and post-coup regime classifications for a panel of 175 countries between 1950 and 2020 (the full coup panel spans 204 countries through 2025, with the 180-country merged sample of Section 4 the active estimation window). The combination (a curated event dataset and an explicit acknowledgment that the events are endogenous to the macro environment) shifted the literature toward identification strategies that account for the reverse channel. Local-projection and event-study designs with pre-event falsification tests (Jordà, 2005; Callaway and Sant'Anna, 2021) became the natural tools for this work.

The third phase, of which the present paper is part, focuses on the *institutional channels* through which coups affect the economy and on the heterogeneity of coup events themselves. The motivating result is the working paper of Bennett et al. (2025), who use labor-productivity data from the Groningen Growth and Development Centre 10-sector database to estimate the effects of coups on productivity in a panel of 39 countries between 1950 and 2012. They find that the productivity effects of coups concentrate in coups that produce a transition from democracy to civilian autocracy, and they articulate the theoretical mechanism in terms of institutional restructuring: changes in property-rights regimes, judicial constraints, and the formal post-coup environment alter the conditions under which firms invest in productivity-enhancing innovation. The BBG framework is distinguished from earlier work in two respects. First, it isolates a specific class of coup events (democracy-destroying transitions) rather than pooling across heterogeneous classes. Second, it identifies a specific channel, formal institutional change, where earlier work invoked generic political uncertainty.

Related strands in the broader political-economy literature support the institutional-channel framing. Acemoglu and Robinson (2008) document the persistence of elite political power across regime changes and the implications for the institutional environment that follows. Meyersson (2014) provides a regression-discontinuity estimate of the economic effects of political-regime transitions that establishes the relevance of within-regime institutional shifts for the macroeconomy.

The earlier literature on the macroeconomic effects of political instability through inflation (Aisen and Veiga, 2007; Cukierman et al., 1992), fiscal policy (Darby et al., 2004), and foreign investment (Chen and Feng, 1996; Sabir et al., 2019) remains relevant for the broader set of outcomes we examine. In the local-projection results of Section 6.1, foreign direct investment, inflation, government spending, and unemployment show no detectable impact-year response in the first-event-only sample. These nulls do not contradict the earlier literature, which used different identification strategies and different operationalizations of instability. They indicate that under the specification we adopt (first-event LP on coup events with country and year fixed effects) the empirical signal for these outcomes does not survive. The paper is therefore most directly positioned as a productivity-channel and institutional-channel contribution, with the broader macro-outcome nulls reported for completeness rather than as the focal claims.

3 Theoretical Framework

3.1 Overview

A single mechanism runs through this section. Political instability is a stochastic hazard, the instantaneous probability of a coup, and every forward-looking agent in the model must act on a future that a coup may cut short. Firms choose when to commit irreversible capital, incumbents choose how much to inflate, households choose whether to emigrate. Each optimises against the same hazard: for the firm it inflates the option value of waiting, for the incumbent it shortens the horizon over which the cost of inflation is borne. The framework is closest to Bennett et al. (2025), who treat regime uncertainty as a distortion of entrepreneurial judgement. The option-value apparatus is that of Dixit and Pindyck (1994).

The model decomposes a coup's policy consequences into two moments. The *mean* of the post-coup policy distribution captures the coalition channel: which group writes the rules after the event determines the level of expropriation, the seigniorage policy, the public-goods bundle. The *variance* captures the uncertainty channel: how dispersed the possible post-coup outcomes are determines the option value of waiting on investment, technology adoption, and hiring. The two channels load differently on different outcomes, and the differential loading generates testable predictions. The full formal derivation is given in Appendix A. Section 7 identifies the structural parameters in two stages. The single-outcome estimation establishes that the uncertainty-channel parameter γ carries the predicted sign while the coalition-channel parameter α is not separately identified by TFP moments alone. The two-block joint estimation across TFP, growth, and investment yields separate identification of α across the two blocks and recovers the cross-block sign heterogeneity the coalition channel exhibits (positive in the productivity block, negative in the investment block), consistent with the within-block versus across-block distinction the option-value derivation supports.

Figure 1 sketches the transmission the rest of this section formalizes: political instability enters as a stochastic coup hazard, a coup realization is summarized by two policy moments, each moment loads on a separate channel of the firm's option-value decision, the firm's adoption-cum-investment decision responds, and it maps to the three macroeconomic outcomes that the empirical section adjudicates.

Figure 1: Transmission from political instability to macroeconomic outcomes

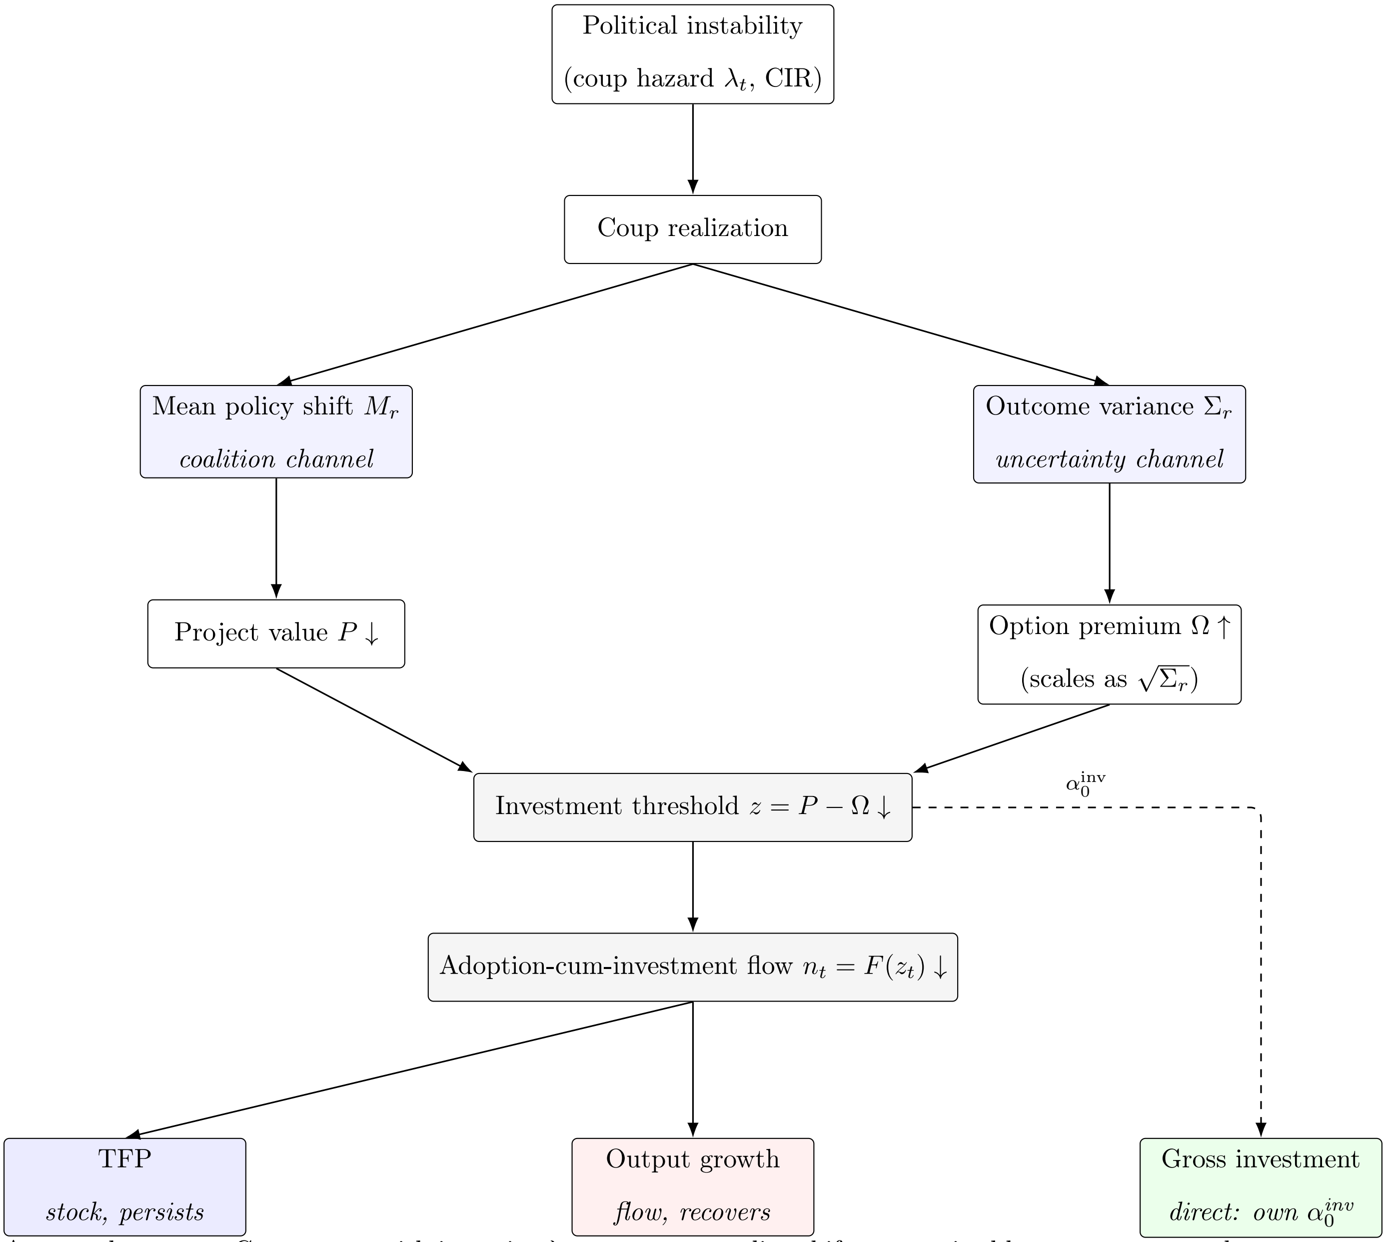

A coup, drawn as a Cox process with intensity λ_t, generates a policy shift summarized by two moments: the mean M_r and the variance Σ_r. The mean loads on the project's expected value through the NPV channel, the variance loads on the option premium of waiting through the standard Black–Scholes $\sqrt{\Sigma_r}$ scaling. Their combination determines the threshold $z = P - \Omega$, which drives the adoption flow $n_t = F(z_t)$. The flow maps to two of the outcomes, TFP as the integral of the adoption flow (stock-like persistence) and output growth as the adoption flow itself (flow-like recovery). Gross investment is instead the firm's direct decision variable, loading on the coalition channel through a separate parameter α_0^{inv} that captures the direct effect of post-coup policy shifts on the project value. The full derivation is in Appendix A.

The framework delivers a set of predictions that the empirical sections adjudicate where the design permits. On *ordering*, the outcome families respond in sequence, economic performance first, macroeconomic stability next, and welfare last. On *persistence*, responses in stock variables are persistent while responses in flow variables are transitory. On the *crossing*, a self-coup behaves like a normal successful coup on variance-driven outcomes and like a failed coup on coalition-driven ones, the signature of the model and the cell that identifies it. On *coup type*, military and civilian coups differ asymmetrically, through the coalition a successful coup installs rather than through the variance it creates. On *institutions*, the per-coup effect scales with ex-ante institutional dispersion, reconciling the political stability of democracies with the severity of a coup when one occurs.

3.2 The entrepreneur and the option-value channel

A coup is, for the people who build factories, a risk to be priced. Forward-looking firms observe the coup hazard λ_t in continuous time and decide *when* to commit irreversible capital. A higher hazard does two things at once. It depresses the net present value of an operating project, because the expected post-coup change in the capital tax is unfavourable. And it raises the option value of waiting, because the dispersion of the possible post-coup tax outcomes makes the bad realisations costly enough that the option of avoiding them gains value.

The relative weight of these two effects determines which outcomes respond at the pre-event horizon ($h = -3, -2$ in the local projections) and which respond only at the event itself. For outcomes whose value depends on the *variance* of post-coup policy (foreign direct investment, total factor productivity, output growth), the option-value channel dominates: a higher hazard reduces investment even before any coup occurs. For outcomes whose value depends on the *mean* post-coup policy (inflation, government spending), the coalition channel dominates, and the response is concentrated at the event.

Proposition 1 (Investment and the coup hazard)**.** *Aggregate investment is strictly decreas-*

ing in the coup hazard λ_t through both the NPV (coalition) and the option-value (uncertainty) channels. Foreign direct investment loads more heavily than domestic investment on both channels, because foreign capital is disproportionately exposed to expropriation.

The full derivation, the CIR specification of the hazard process, the regime-specific calibration to the Bjørnskov–Rode panel, and the formal statement of Proposition 1 are given in Appendix A.

The framework therefore predicts that investment and FDI decline as the coup hazard rises, the decline appears before the coup, because forward-looking firms respond to λ_t rather than to the realised event, and it is largest for FDI. Investment moves first, placing the performance family at the head of the transmission ordering.

3.3 Total factor productivity and persistence

Productivity rises when firms adopt better technology. Adoption, like investment, is irreversible and carries a sunk cost, so it is governed by the same option logic: a firm adopts when the value of the new technology, net of the option of waiting, clears the sunk cost. The adoption flow x_t is therefore decreasing in the coup hazard, $x'(\lambda) < 0$, by Proposition 1 applied to the technology-choice margin.

The difference between productivity and investment is not in the decision but in the *accumulation*. Investment and FDI are flows: when λ reverts after a coup, the flow returns to its normal level. Total factor productivity is a stock, the accumulated body of adopted technology,

$$\ln A_t = \ln A_0 + \int_0^t x(\lambda_s)\, ds. \tag{1}$$

A coup episode depresses the adoption flow for the duration of the high-λ window. When λ reverts, x recovers, but $\ln A_t$ does not: it remains permanently below its counterfactual path by the cumulative shortfall $\int(\bar{x} - x(\lambda_s))\, ds$ accrued during the episode.

This is the model's persistence prediction. The performance flows (investment, FDI, output growth) trace V-shaped responses that return toward trend once the hazard subsides, while total factor productivity, the integral of a disrupted flow, traces a downward step that does not return. The direct test of this prediction is the cumulative effect comparison reported in Section B.3 (Table 10): under perfect stock persistence the cumulative TFP coefficient at horizon $H = 5$ should approach $(H+1)\beta^A_{c,0} = 6 \cdot \beta^A_{c,0}$, while a fully recovered flow yields $B^y_{c,5} \approx \beta^y_{c,0}$. The empirical amplification ratios at the DEM $\rightarrow$ AUT cell are 5.57 for TFP and 0.52 for growth, consistent with the stock-flow distinction. The F2-timing comparison in Appendix F.1 provides a secondary check.

3.4 The incumbent and the coup-branch predictions

Two further mechanisms complete the framework: a level-channel inflation response we set out but do not adopt empirically, and a branch-crossing structure underpowered in the first-event-only sample. First, the incumbent's seigniorage problem generates the *level-channel* response: because the coup hazard enters the incumbent's effective discount rate (Cukierman et al., 1992), inflation responds to the mean policy shift (the coalition channel) rather than to outcome dispersion (the variance channel), a level-versus-variance split that motivates the cross-block coalition-versus-uncertainty distinction the structural section develops. Second, the three coup branches, failed (F), self-coup (SC), and normal successful (N), generate a crossing structure in which the self-coup is the cell that identifies the crossing, behaving like a normal coup for variance-driven outcomes (TFP, FDI) and like a failed coup for level-driven outcomes (inflation, government spending). The military or civilian character of the installed coalition then refines the normal-successful branch, with military coups predicted to depress long-run productivity by more than civilian ones (Bennett et al., 2021). Because the self-coup is the identifying cell for the crossing and the first-event-only sample contains only four self-coup events (Section 4.2), the crossing cannot be decisively tested, and the thin military cell leaves the coup-type asymmetry similarly underpowered, so both are illustrated rather than adjudicated (Section 3.6). The full derivations (Proposition 2 and the inflation level-signature, the branch prediction table, and the crossing diagram) are in Appendix A.

3.5 Institutions as a shock absorber

The cost of a coup is not the same everywhere, and the model says why. A coup's damage is the option-value response of Section 3.2, and two regime features scale it. The first is surprise. A coup is measured against its pre-event baseline, where the long-run hazard is low, that baseline carries little priced-in coup risk, so a coup is a large and largely unanticipated shock. Where the long-run hazard is high, coup risk is chronically priced in, the baseline is already depressed, and the same coup moves the economy less. The second feature is the variance of post-coup policy: the more dispersed the outcomes a coup can bring, the larger the option-value response.

The framework also predicts how institutions absorb the shock. Separate the macroeconomic burden of political instability into a per-coup effect and a frequency. Strong institutions deliver a *large* per-coup effect, because a coup is a rare surprise, but a *low* frequency of coups. The aggregate burden is the product of the two, and it is small. Weak institutions deliver a muted per-coup effect but a high frequency, and a large aggregate burden. The aggregate cost of political instability is thus lower where institutions are stronger, not because each coup is milder there, but because coups are rare. This resolves an apparent paradox: a parliamentary democracy is politically stable, yet a coup in one is unusually damaging.

Empirical proxy for the variance margin. The variance margin is not directly observable, but the BBG-style transition decomposition of Section 6.2 provides an empirical proxy: the post-coup regime trajectory captures the institutional distance a coup traverses, with democracy-destroying transitions on one end and intra-democracy outcomes (failed coups and no-change successful coups) on the other. The decomposition is the empirical counterpart of the institutions prediction employed in this paper.[1]

[1] A full structural estimation of the variance margin from the Bjørnskov–Rode transition matrix is left to future work.

3.6 Empirical scope

The predictions stated above are theoretical implications of a single regime-uncertainty mechanism. The empirical design in Section 6 discriminates among them where the discrimination is feasible, but it does not adjudicate all five with equal force. The distinction matters for the interpretation of the evidence we report, and we are explicit about it.

Two of the predictions admit empirical tests within the local-projection and joint structural specifications we estimate with the inferential force of the test. The ordering prediction is adjudicated by the two-block structural identification in Section B.2, which recovers the option-value channel parameters with the predicted signs in the productivity block ($\hat{\alpha}_0 > 0$, $\hat{\gamma}_0 < 0$) and the investment block ($\hat{\alpha}_0^{\text{inv}} < 0$, $\hat{\gamma}_0^{\text{inv}} < 0$ at the point estimate). The uncertainty channel carries the predicted negative sign in both blocks across all empirically plausible calibrations of the dispersion moment. Under the primary country-cluster bootstrap that propagates moment uncertainty, $\hat{\gamma}_0$ is not separately distinguishable from zero, so we read the structural estimates as a sign-and-bound decomposition rather than point identification (Section B.2). The coalition channel exhibits the cross-block sign heterogeneity the theory predicts in the point estimates at the headline calibration (positive in the productivity block, negative in the investment block), though only the investment-block loading is separately distinguishable from zero (the productivity-block loading's interval crosses zero and its sign reverses under one alternative V-Dem calibration, Appendix G.3), consistent with the within-block versus across-block distinction articulated in Appendix A.2.3. The persistence prediction is tested directly by the cumulative effect comparison in Section B.3 (Table 10): the TFP cumulative coefficient at $H = 5$ delivers an amplification ratio of 5.57 (near the perfect-persistence value of $H + 1 = 6$) with cumulative loss -0.236 at $p < 0.001$ under the independence approximation (Table 10), while the growth amplification is 0.52 with cumulative response not statistically distinguishable from zero, the TFP–growth asymmetry under the F2-timing alternative of Appendix F.1 provides a secondary check. Three predictions are illustrated by the data without rising to statistically powered adjudication.

The institutions prediction is partially illustrated by the BBG-style transition decomposition (Section 6.2), which provides a direct empirical proxy for the institutional-variance margin of the model: the impact-year productivity contraction concentrates in democracy-destroying coups (-0.042, $p < 0.001$), is statistically zero in democracy-preserving coups (-0.012), and sits at intermediate magnitudes for intra-autocratic transitions, an ordering consistent with the prediction that the per-coup effect scales with the institutional distance the coup traverses. The aggregate burden-of-instability statement the prediction also implies (that the per-coup effect times the frequency of coups yields a small aggregate cost in strong-institutional countries and a large aggregate cost in weak-institutional countries) is not estimated and is left to future work. The self-coup crossing is not decisively testable under the first-event-only sample restriction we adopt for identification, which reduces the self-coup cell to four events, below the power required to distinguish the self-coup coefficient from the failed and normal-successful coefficients with statistical confidence. The military–civilian asymmetry is similarly underpowered: the pre-coup regime heterogeneity reported in Appendix F.2 does not reject joint equality of effects across pre-coup regime classes at the 10% level, and the military category is itself thin in first-event-only samples ($n = 9$) for the substantive reason discussed in Section 4.2, military dictatorships typically emerge from coups rather than precede them. The military–civilian asymmetry the theory predicts is not empirically distinguished from null by the design.

The mapping is the empirical scope of the paper. The framework is the theoretical scaffolding within which we interpret the evidence, the design adjudicates a subset of those predictions. Broader theoretical implications of the regime-uncertainty mechanism that are beyond the inferential reach of our specification (elite-bargaining adjustments after regime change, the path-dependence of post-coup institutions, the ideological orientation of post-coup policy) are plausible within the framework but are not tested here.

4 Data

We construct a merged country-year panel covering 180 countries from 1960 to 2022, the estimation window for the local-projection analysis below, bounded by PWT 11.0 outcome coverage. The underlying Bjørnskov–Rode coup panel covers 204 countries through 2025 and is summarized over its full range in Table 2, the merged 180-country panel is the active sample for outcome-based estimation. We draw on the World Bank's World Development Indicators, the Penn World Table version 11.0 (Feenstra et al., 2015), and the Bjørnskov–Rode Regime Dataset. The outcome variables measure macroeconomic and welfare performance, the treatment variable is the incidence of coups d'état.

4.1 Dependent Variables

The analysis focuses on total factor productivity as the primary outcome of theoretical interest, with output growth as a closely related secondary outcome. Six additional macroeconomic outcomes are examined for completeness and reported in tables and the appendix: foreign direct investment, trade openness, inflation, government spending, unemployment, and net migration. The empirical results in Section 6 concentrate on the first two outcomes, the remaining six are reported as nulls under the identification strategy we adopt and are not the focal claims of the paper.

Total factor productivity is the `rtfpna` index from the Penn World Table 11.0 release (Feenstra et al., 2015), the most recent revision of the dataset incorporating updated price benchmarks and capital stock corrections. The series enters as ln(rtfpna), so the LP coefficient $\hat{\beta}_h$ is a log change, approximately a proportional change for small effects (for example, $\hat{\beta}_0 = -0.042 \approx -4.2\%$). The focal-cell first-event estimates are robust to using earlier PWT releases. The PWT 11.0 release does affect the multi-period staggered diagnostic of Appendix H: pre-1972 TFP coverage for several Latin American countries available only under PWT 11.0 brings Argentina 1966 (Onganía) into the set of dCDH contributing switchers, where it was previously excluded by missing lagged coverage in the earlier release. Output growth is the annual rate of real GDP growth. FDI and trade openness measure external linkages, both expressed in nominal terms. Government spending is expressed as a share of GDP. Inflation is the first difference of the log consumer price index, $\Delta \ln(\mathrm{CPI}_{it}) \times 100$. Unemployment is the World Bank's unemployment rate. Net migration captures the difference between immigration and emigration (Gebremedhin and Mavisakalyan, 2013).[2]

To make effects comparable across outcomes, we place all dependent variables on a logarithmic scale using a consistent criterion. Rate and ratio variables (growth, unemployment, government spending) enter as $\ln(1 + x/100) \times 100$, a transformation that is well defined at zero and interpretable in percentage points. Strictly positive variables (trade openness, TFP) enter in natural logarithms. Variables that can take non-positive values (FDI and net migration) are transformed following Chen and Roth (2023), adding the absolute value of the sample minimum plus a small constant before taking logs. Inflation is measured as the first difference of the log consumer price index. All dependent variables are winsorized at the 1st and 99th percentiles.

4.2 Coups d'État as Events of Political Instability

We measure political instability with coups d'état, drawing on the Bjørnskov–Rode Regime Dataset, version 6.2 (Bjørnskov and Rode, 2020). A coup is a discrete, dated, and externally observable attempt by political or military insiders to remove the sitting executive. This is a deliberate departure from perception-based indices of instability, whose content is difficult to separate from general institutional quality (Langbein and Knack, 2010): a coup either occurs in a given country-year or it does not, and its timing is recorded rather than inferred. The dataset records 569 coup attempts in 114 countries between 1950 and 2025. After the Powell–Thyne re-timing reconciliation (footnote 4 below shifts two coups originally dated 1950 to 1949, before the start of the BR panel proper), the constructed coup panel contains

[2]All variables except TFP are from the World Bank's World Development Indicators.

567 events (Table 2), of which 310 fail and 257 succeed. The mode breakdown on the 569 raw count is 414 military, 143 civilian, and 12 royal. The two re-timed coups are excluded from the constructed panel and therefore from the 257 / 310 success-failure totals in Table 2. Merged into the country-year panel by ISO country code, coups become the treatment whose macroeconomic consequences the remainder of the paper traces.

Rather than treating every coup alike, we separate coups into three arms entered as parallel indicators with non-coup country-years as the omitted category. *Failed coups* (298 country-years) are attempts that are defeated, leaving the incumbent regime in place, a failed coup realizes the uncertainty of an openly contested executive without producing any change in leadership or formal institutions, and so isolates the pure uncertainty channel. *Self-coups* (17) are successful coups in which the sitting head of state remains in office and consolidates power, the classic *autogolpe*, uncertainty is resolved through entrenchment rather than replacement. *Normal successful coups* (230) are successful coups that install a new leader or ruling coalition, where the uncertainty shock is compounded by an actual transfer of power.

The three-arm decomposition is reported in Section 6 as background structure on the data. The main empirical specification (Section 5) uses a four-way BBG-style transition decomposition by the pair (pre-coup regime, post-coup regime), which is the empirical operationalization of the institutional-replacement channel central to the paper. Twenty country-years contain coups of more than one arm, a robustness check drops them.

A self-coup is identified mechanically. A successful coup is coded as a self-coup when the Bjørnskov–Rode head-of-state record names the same individual as the executive in the year before, the year of, and the year after the event. The three-year window distinguishes a genuine *autogolpe* from a coup followed by the rapid restoration of the prior incumbent.[3]

[3]The seventeen self-coups are: the Democratic Republic of the Congo (1961), Nigeria (1967), Syria (1969), Turkey (1971), Burundi (1972), Uruguay (1973), Cyprus (1974), Maldives (1975), Ghana (1976), Madagascar (1976), Bangladesh (1977), Fiji (1988), Peru (1992), São Tomé and Príncipe (2003), Fiji (2006), Tunisia (2021), and Sudan (2022). Peru's 1992 *autogolpe* under Alberto Fujimori is the motivating case. The strict three-year rule excludes coup-plus-restoration episodes such as Sierra Leone (1997).

Because the local-projection design in Section 5 dates every response relative to the event year, the year assigned to each coup is consequential. We therefore cross-check the Bjørnskov–Rode datings against an independent catalogue of coups, that of Powell and Thyne (2011). Of the 526 Bjørnskov–Rode coup country-years, 439 match a Powell–Thyne coup within one year: 290 are dated to the same year, 131 are dated one year later by Bjørnskov–Rode, and 18 one year earlier. The remaining 87 have no Powell–Thyne counterpart within a year, and 43 differ by two years or more. The asymmetry of the one-year cases (131 late against 18 early) indicates a systematic tendency in the integrated file to assign certain coups to a year adjacent to the one in the source catalogue. For the 137 coup events involved we adopt the Powell–Thyne year, the 18 events that Bjørnskov–Rode dates one year earlier are left unchanged, a conservative choice that avoids manufacturing agreement. All results use this reconciled timeline.[4]

An important descriptive implication of the first-event-only restriction adopted in Section 5 is that military dictatorships become comparatively rare pre-coup states. In the full Bjørnskov–Rode coup panel, military regimes account for 38.7% of coup-year cells (203 of 525), in the first-event-only sample, only 9 of the 110 first-coup events occur in a country whose regime is classified as military the year before. This compositional shift suggests that many military autocracies in repeated-coup samples are themselves endogenous products of prior coups, with the repeated-coup cycles in countries such as Bolivia, Chad, Sudan, and Burkina Faso emerging from a first coup against a democracy or civilian dictatorship that installs a military regime which then experiences subsequent military-on-military coups.

Coups are the measure of political instability throughout the main analysis.

[4]Re-timing shifts two coups originally dated 1950 to 1949, before the start of the Bjørnskov–Rode panel, the constructed coup panel therefore contains 567 of the 569 recorded coup events. Both predate the 1960–2022 estimation window and affect no estimate. The 567 / 310 failed / 257 successful totals in Table 2 reflect the post-re-timing panel. The 569 / 310 / 259 / 414 / 143 / 12 raw counts reported in the text above refer to the unreduced BR record.

4.3 Control Variables

We include three blocks of control variables. The first controls for initial economic conditions through lagged GDP per capita, following Acemoglu et al. (2008) and Aisen and Veiga (2013). The second captures demographic and human capital characteristics: years of schooling (Barro and Lee, 2010), total population, and the urbanization rate. The third proxies for physical capital and infrastructure: natural resource rents, gross fixed capital formation, and R&D expenditure. Section 5 introduces these controls incrementally. R&D expenditure is available for far fewer country-years, so we report it in a separate robustness column rather than in the preferred specification.

4.4 Descriptive Statistics

Table 1 reports summary statistics for the outcome and control variables, grouped by family. Economic growth averages 3.6 percent per year. The welfare variables have smaller samples, about 4,000 country-year observations, because they are interpolated only within the survey range. All dependent variables are winsorized at the 1st and 99th percentiles. Table 2 summarizes the coup treatment: its distribution across decades, and the split between failed and successful events. Coups are concentrated in Sub-Saharan Africa and Latin America but appear in every region, and they are spread across the sample period, with a peak in the 1960s and 1970s and a gradual decline thereafter.[5]

[5] The descriptive count of 567 coup events refers to the Bjørnskov–Rode panel after the Powell–Thyne re-timing reconciliation (569 raw events minus two re-timed to 1949 and therefore before the estimation window). The three-arm decomposition (298 failed, 17 self-coups, 230 normal successful) sums to 545 events because the remaining 22 are coup-year cells that the BR coding flags as transitional or as multi-arm country-years (the twenty multi-arm country-years noted earlier in this subsection) and are absorbed into the overall coup indicator rather than classified into a single arm. The first-event-only restriction adopted in Section 5 further reduces the sample to 110 first-coup events across 110 countries (one per country). All enter the LP estimation, with the number contributing to a given outcome at $h = 0$ varying with covariate coverage.

Table 1: Descriptive Statistics

Variable	Obs.	Mean	Std. Dev.	Min	Max
Panel A: Macroeconomic Performance					
Economic Growth (%)	10,096	3.61	5.19	-14.14	20.80
Total Factor Productivity	6,697	1.05	0.42	0.35	3.45
Foreign Direct Investment (mn USD)	9,136	4,298	14,697	$-2,138$	105,293
Trade (% of GDP)	8,637	77.77	49.87	12.42	316.48
Trade Openness (mn USD)	7,317	115,578	299,353	57	1,844,136
Net Migration (millions)	13,888	-0.00	0.08	-0.37	0.44
Panel B: Macroeconomic Stability					
Inflation Rate (%)	8,506	9.25	22.98	-19.41	477.49
Unemployment Rate (%)	6,168	8.17	5.92	0.60	27.70
Government Spending (% GDP)	8,295	16.15	6.94	4.84	45.78
Panel C: Coup Treatment Indicators					
Any coup	12,852	0.0358	0.186	0	1
Failed coup	12,852	0.0202	0.141	0	1
Self-coup (*autogolpe*)	12,852	0.0013	0.036	0	1
Normal successful coup	12,852	0.0156	0.124	0	1
Military coup	12,852	0.0257	0.158	0	1
Civilian coup	12,852	0.0106	0.102	0	1
Panel D: Control Variables					
GDP Per Capita (USD)	10,337	11,029	12,337	713	59,120
Years of Schooling	7,042	6.48	3.37	0.04	13.74
Population (millions)	13,641	24.89	104.30	0.003	1,420
Urban Population (%)	13,545	51.63	25.71	2.08	100
Gross Fixed Capital Formation (% GDP)	7,720	22.20	8.11	-2.42	93.55
R&D Expenditure (% GDP)	2,237	0.95	0.96	0.01	5.71
Natural Resource Rents (% GDP)	9,235	6.83	11.05	0	88.59

Notes: This table reports summary statistics for the variables used in the analysis, organized by outcome family. All dependent variables are winsorized at the 1st and 99th percentiles. Coup treatment indicators (Panel C) are observed for the full merged country-year panel of 12,852 observations covering 1960–2022. The difference in N between Panel C and the outcome panels reflects missing data in the outcome variables. The local-projection estimation samples in Sections 6 and 6.4 apply additional restrictions from the first-event identification design and are smaller than the descriptive coverage reported here. Negative FDI values reflect net disinvestment. For variables that can take non-positive values (FDI and net migration), the log transformation follows Chen and Roth (2023), adding the absolute value of the minimum plus 0.001 before taking logs.

The full coup dataset spans 1950 to 2025 and contains 567 events. The decade distribution is reported in Table 2. Coups peak in the 1960s and 1970s and decline thereafter, with the regional distribution weighted toward Sub-Saharan Africa and Latin America. The estimation sample in the local-projection analysis below covers 1960 to 2022, the 1950s row is reported for completeness only.

Table 2: Coups d'État by Decade, 1950–2025

Decade	Failed	Successful	Total
1950s	30	28	58
1960s	50	61	111
1970s	55	62	117
1980s	65	40	105
1990s	54	28	82
2000s	27	18	45
2010s	16	10	26
2020s	13	10	23
Total	310	257	567

Coup events from the Bjørnskov–Rode Regime Dataset v6.2, dated on the timeline reconciled with Powell and Thyne (2011). The dataset records 569 coups in total, re-timing shifts two coups originally dated 1950 to 1949, before the start of the panel, so the constructed coup panel contains 567 events. Of the 569 recorded coups, Sub-Saharan Africa accounts for 243, Latin America and the Caribbean for 127, the Middle East and North Africa for 67, East Asia and the Pacific for 51, South Asia for 27, and Europe and Central Asia for 23, the remaining 31 are in small states and territories not classified to a region.

5 Empirical Strategy

The empirical strategy has three components: a local-projection specification with country and year fixed effects estimated on a sample restricted to the first coup observed in each country, a decomposition of those projections by the regime trajectory the coup produces, following Bennett et al. (2025), and a robustness architecture concentrated on the focal cell that the decomposition identifies. The strategy is associational, we do not impose strict parallel-trends assumptions as identification restrictions, and we describe what the robustness exercises do and do not establish in Section 5.5.

5.1 Local-projection specification

For each outcome Y_{it} and horizon $h \in \{-3, -2, 0, +1, \ldots, +8\}$ we estimate the regression

$$Y_{i,t+h} - Y_{i,t-1} = \alpha_i^h + \gamma_t^h + \beta_h \cdot \text{shock}_{it} + \phi^h Y_{i,t-1} + \psi^h \ln \text{GDP}_{i,t-1} + \varepsilon_{i,t+h}^h, \tag{2}$$

where shock_{it} is an indicator equal to one in the year of the first coup observed in country i and zero elsewhere. The coefficient β_h estimates the average conditional dynamic response to a first coup, evaluated at horizon h, with the year preceding the event ($h = -1$) implicitly serving as the reference. Country fixed effects α_i^h absorb time-invariant level differences across countries, year fixed effects γ_t^h absorb common shocks. The lagged outcome and lagged log GDP per capita are included as controls, standard errors are clustered at the country level. We estimate (2) separately at each horizon, following Jordà (2005) and the more recent applied discussion in Plagborg-Møller and Wolf (2021).

The pre-event horizons $h = -3$ and $h = -2$ are included as falsification. Under the conditional-associations framing we adopt, the pre-event coefficients are part of the empirical pattern rather than an identifying restriction, we report them, test them for nullity, and acknowledge the implication for interpretation. The horizon $h = +5$ in the post-event window is reported as a summary of medium-run persistence, horizons $h = +6$ through $h = +8$ are reported in the appendix and used to verify that the persistence of the focal cell does not reverse over the longer window.

5.2 First-event restriction

The estimation sample is restricted to country-years prior to the second coup in country i. The restriction has two parts. First, the shock$_{it}$ indicator equals one only in the year of the first coup observed for country i, subsequent coups in the same country do not enter the regression as additional shocks. Second, the country's panel is truncated at the year of its second coup, so that no country-year observation contributes to the regression after the second coup has occurred. The reason for both restrictions is the same: in repeated-coup countries, the post-event horizons of the first coup overlap with the pre-event horizons of subsequent coups, contaminating the estimated β_h. This is the same first-event identification used in the staggered difference-in-differences literature (Callaway and Sant'Anna, 2021).

Beyond the econometric motivation, the first-event design also matches the theoretical object of interest: the option-value framework of Section 3 models the firm's response to the *first* major institutional shock under uncertainty, not the marginal effect of an additional shock in a cycling-regime environment. The choice of first-event LP as the primary specification is therefore not a statistical convenience but a correspondence with the structural model. A multi-period staggered comparison reported in Appendix H confirms the magnitude under both absorbing (Callaway–Sant'Anna) and non-absorbing (de Chaisemartin–D'Haultfœuille) designs.

The restriction has substantive consequences for the sample. The Bjørnskov–Rode panel contains 525 coup-year cells across 114 countries, the first-event-only sample contains 110 first-coup events across 110 countries. All enter the estimation, with the number contributing to a given outcome varying with covariate coverage. The truncation at the second coup reduces the country-year sample from 13,055 to 9,137 observations. The compositional shift in pre-coup regime types implied by the restriction is discussed in Section 4.2, in particular, military dictatorships, which account for 38.7% of coup-year cells in the full panel, contribute only 9 observations to the first-event-only sample, reflecting the historical regularity that military regimes are typically the consequence of an earlier coup rather than a stable starting state. By Bjørnskov–Rode coup type, the 110 first coups comprise 53 successful coups that produce a normal post-coup government, 59 failed attempts, and 4 self-coups by the incumbent executive. These type counts sum to 116 rather than 110 because six first-coup years contain more than one classified coup. The outcome variables are total factor productivity, real output growth, foreign direct investment, trade openness, inflation, government spending, unemployment, and net migration, all in logs (inflation in log first differences) and winsorized at the 1% level.

5.3 Transition decomposition

For the decomposition in Sections 6.2 and 6.4, we classify each first coup by the pair (pre-coup regime, post-coup regime), using the Bjørnskov–Rode regime classification. The pre-coup regime is measured at $t-1$, the year preceding the coup. The post-coup regime is measured at $t+1$ with a fallback to $t+2$ when the $t+1$ value is missing or transitional. We group the pre-coup and post-coup classifications into democratic (parliamentary, mixed, and presidential democracies, codes $\{0,1,2\}$) and autocratic (civilian, military, and royal dictatorships, codes $\{3,4,5\}$) and define four transition categories. DEM $\to$ AUT comprises events with pre-coup democracy and post-coup autocracy ($n = 14$ in the first-event-only sample), AUT $\to$ DEM comprises pre-coup autocracy and post-coup democracy ($n = 2$), INTRA_AUT comprises events that are autocratic both before and after ($n = 65$), and INTRA_DEM comprises events that remain democratic both before and after ($n = 18$), containing 16 failed coups and 2 successful coups that did not produce a regime change and functioning as the empirical counterfactual to the DEM $\to$ AUT class in Section 6.3. Regime classifications outside $\{0,\dots,5\}$ (transition states, occupation regimes), as well as observations with missing pre-coup or post-coup regime, are absorbed by a separate OTHER indicator ($n = 11$) in the decomposition and not reported in the focal tables. The four-class breakdown therefore covers $14+2+65+18 = 99$ events, and the OTHER category absorbs the remaining 11 first-coup events, summing to the 110-event LP-contributing sample.

The decomposed local projection replaces the single shock_{it} indicator in (2) with four transition-class indicators entered jointly:

$$
\begin{aligned}
Y_{i,t+h} - Y_{i,t-1} = \alpha_i^h + \gamma_t^h & \\
& + \beta_h^{DA}\, \mathbb{1}[\text{DEM} \to \text{AUT}]_{it} + \beta_h^{AD}\, \mathbb{1}[\text{AUT} \to \text{DEM}]_{it} \\
& + \beta_h^{AA}\, \mathbb{1}[\text{INTRA\_AUT}]_{it} + \beta_h^{DD}\, \mathbb{1}[\text{INTRA\_DEM}]_{it} \\
& + \phi^h\, Y_{i,t-1} + \psi^h \ln \text{GDP}_{i,t-1} + \varepsilon_{i,t+h}^h .
\end{aligned} \tag{3}
$$

The four coefficients are estimated jointly within the same regression, which preserves the comparison of magnitudes across classes and permits Wald tests of pairwise equality. The empirical counterfactual underlying the identification logic of Section 6.3 is the pairwise contrast $\beta_h^{DA} - \beta_h^{DD}$: the difference between democracy-destroying coups and democracy-preserving coup attempts in the same outcome at the same horizon.

Using democracy-preserving coup attempts as the counterfactual for regime-destroying ones is in the spirit of Jones and Olken (2009), who contrast successful and failed assassination attempts on national leaders. We read $\beta_h^{DA} - \beta_h^{DD}$ as an associational contrast conditional on the controls in (3) rather than as the effect of an as-good-as-random regime change.

5.4 Robustness architecture

The robustness exercises in Section 6.5 are concentrated on the DEM $\to$ AUT TFP coefficient at $h = 0$, the focal cell that the decomposition identifies. Four checks are reported. The first is a *country jackknife*: the decomposition (3) is re-estimated 14 times, dropping each of the 14 DEM $\to$ AUT countries one at a time, to test whether any single country drives the focal estimate. The second is *leave-one-region-out*: the decomposition is re-estimated seven times, dropping each World Bank region in turn, with particular attention to Latin America and the Caribbean, which accounts for 50 percent of the DEM $\to$ AUT events. The third is the *alternative regime-classification timing*: the decomposition is re-estimated with the post-coup regime measured at $t+2$ rather than $t+1$, testing whether the focal estimate depends on the particular year used to classify the post-coup regime and distinguishing transient classification artifacts from a durable institutional effect. The fourth is the *joint pre-event nullity test*: a stacked specification with horizon-interacted fixed effects estimates β_{-3}^{DA} and β_{-2}^{DA} jointly and tests their joint nullity, assessing whether the focal estimate reflects a response to the event or the continuation of a pre-existing trajectory.

The robustness battery is targeted at the focal cell rather than distributed across all decomposition cells. This choice reflects the tradeoff between thoroughness and multiple-comparisons inflation in a sample with $n = 110$ first-coup events and four decomposition classes, of which the focal cell has 14 events.

5.5 Identification framing

We adopt a conditional-associations framing throughout. The local projections (2) and (3) estimate the response of each outcome to a first coup, conditional on country and year fixed effects, lagged outcome, and lagged log GDP per capita. We do not impose strict parallel-trends assumptions as identification restrictions, and we do not have a quasi-experimental source of exogenous variation in coup occurrence. Where pre-event coefficients are non-zero we report them, and where the joint pre-event nullity test rejects we describe the implication for the interpretation of the focal estimate (Section 6.5).

The empirical content of the strategy is that of the four robustness exercises in Section 5.4, taken together. We do not claim that the pattern identified by the decomposition is a structural causal parameter free of selection. We claim that the pattern survives the natural robustness exercises and is most consistent with an interpretation in which the dispersion of post-coup institutional outcomes (formalized in Section 3 as the variance margin of the option-value framework of Dixit and Pindyck (1994)) is the operative pathway through which coups affect productivity, rather than the coup event itself or generic political uncertainty alone. The identification logic parallels that of Bennett et al. (2025) applied to the Groningen 10-sector labor-productivity panel.

6 Results

The empirical analysis proceeds in three stages of increasing specificity. We first estimate aggregate dynamic responses to the first coup observed in each country, using local projections with country and year fixed effects and standard errors clustered by country. We then decompose those responses by the regime trajectory the coup produces, following the framework of Bennett et al. (2025). We finally concentrate on the focal cell that the decomposition identifies (coups that move a country from democracy to autocracy) and submit it to four robustness exercises designed to discipline the inference. The structure of the section mirrors the progressive narrowing of the identification claim.

6.1 Aggregate dynamic responses

Table 3 reports $\hat{\beta}_h$ at the seven focal horizons. Two outcomes show statistically meaningful impact-year contractions: total factor productivity ($\hat{\beta}_0 = -0.020$, $p < 0.05$) and output growth ($\hat{\beta}_0 = -2.23$, $p < 0.01$, expressed in percentage points). The TFP contraction extends to $h = +1$ ($\hat{\beta}_1 = -0.025$, $p < 0.10$) and fades by $h = +2$. The growth response is concentrated at the impact year and does not persist materially: $\hat{\beta}_h$ is statistically indistinguishable from zero at every post-impact horizon in the aggregate specification. Net migration is positive at $h = +5$ ($\hat{\beta}_5 = +0.19$, $p < 0.05$), foreign direct investment, trade openness, inflation, government spending, and unemployment show no statistically detectable impact at any horizon.

Figure 2 plots the TFP and growth responses with 95 percent confidence bands. The two outcomes that move at the impact year do so by economically substantive magnitudes (approximately a 2 percent TFP contraction and a 2.2 percentage-point output-growth contraction) but the aggregate estimate averages across heterogeneous coup events of distinct institutional character. The decomposition in the next subsection unpacks that heterogeneity.

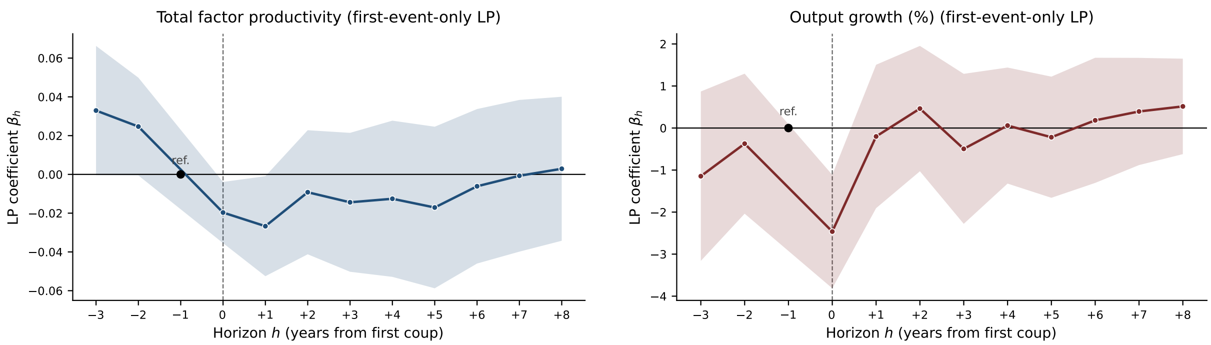

Figure 2: Aggregate local-projection responses to the first coup: total factor productivity (left) and output growth (right). First-event-only sample, truncated at the second coup. Shaded bands are 95% confidence intervals from standard errors clustered at the country level. The reference point at $h = -1$ is normalized to zero by the LP specification.

Table 3: Aggregate local projections, first-event-only sample

Outcome	$h=-3$	$h=-2$	$h=+0$	$h=+1$	$h=+2$	$h=+3$	$h=+5$
TFP	$+0.033^{*}$	$+0.025^{*}$	-0.020^{**}	-0.027^{**}	-0.009	-0.014	-0.017
	(0.017)	(0.013)	(0.008)	(0.013)	(0.016)	(0.018)	(0.021)
Growth	-1.146	-0.372	-2.463^{***}	-0.201	$+0.462$	-0.495	-0.219
	(1.028)	(0.850)	(0.689)	(0.870)	(0.761)	(0.910)	(0.735)
FDI	$+0.076$	-0.038	-0.020	$+0.028$	$+0.053$	$+0.101$	$+0.105$
	(0.064)	(0.071)	(0.041)	(0.056)	(0.064)	(0.071)	(0.090)
Trade openness	$+0.039$	$+0.031$	-0.041	$+0.045$	$+0.098$	$+0.046$	-0.060
	(0.031)	(0.025)	(0.027)	(0.059)	(0.067)	(0.064)	(0.071)
Inflation	-0.277	-1.349	$+3.866$	-0.662	-1.164	-0.910	-1.284
	(1.385)	(1.414)	(2.636)	(2.328)	(1.808)	(3.383)	(1.795)
Gov. spending	-0.023	$+0.163$	-0.455^{*}	-0.192	$+0.135$	$+0.700$	-0.067
	(0.273)	(0.221)	(0.253)	(0.457)	(0.405)	(0.505)	(0.441)
Unemployment	-0.494	-0.046	-0.117	-0.255	$+0.020$	-0.168	$+0.215$
	(0.642)	(0.279)	(0.241)	(0.367)	(0.384)	(0.563)	(0.640)
Net migration	-0.020	-0.012	-0.188	-0.169	-0.098	-0.047	$+0.198^{**}$
	(0.028)	(0.020)	(0.202)	(0.182)	(0.169)	(0.156)	(0.097)

Notes. Coefficients on the first-coup indicator from horizon-h local projections of the form $Y_{i,t+h} - Y_{i,t-1} = \alpha_i^h + \gamma_t^h + \beta_h \cdot \text{coup}_{it}^{\text{first}} + \phi^h L.Y_{it} + \psi^h L.\ln\text{GDP}_{it} + \varepsilon_{i,t+h}^h$. Sample restricted to country-years prior to the second coup. $h{=}-3, -2$ are pre-event falsification horizons under the conventional causal reading, or part of the descriptive trajectory under the associations framing adopted here. Standard errors clustered by country. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.2 Transition decomposition

Bennett et al. (2025) argue that the productivity-relevant consequences of coups operate primarily through institutional change rather than through the coup event itself. To assess this, we classify each first coup by the pair (pre-coup regime, post-coup regime), measured one year before and one year after the coup using the Bjørnskov–Rode regime classification. Three of the four resulting transition categories are informative, the fourth (autocracy-to-democracy) has only two events and is reported for completeness only.

Table 4: BBG-style transition decomposition: local projections by post-coup regime trajectory

	$h=-3$	$h=-2$	$h=+0$	$h=+1$	$h=+5$
TFP					
DEM→AUT (n=14)	+0.012	+0.025*	−0.042***	−0.030**	−0.030***
	(0.024)	(0.015)	(0.011)	(0.013)	(0.011)
INTRA_AUT (n=65)	+0.051**	+0.038*	−0.016	−0.037*	−0.006
	(0.026)	(0.021)	(0.012)	(0.020)	(0.036)
INTRA_DEM (n=18)	+0.004	+0.003	−0.011	+0.010	−0.016
	(0.010)	(0.015)	(0.010)	(0.015)	(0.017)
AUT→DEM (n=2)	−0.016	−0.024***	−0.042	−0.083	−0.108**
	(0.015)	(0.004)	(0.032)	(0.058)	(0.046)
Growth					
DEM→AUT (n=14)	+0.270	+1.907	−3.689**	+2.086	−0.674
	(1.348)	(1.174)	(1.485)	(1.516)	(1.316)
INTRA_AUT (n=65)	−1.421	−0.408	−2.522***	−0.998	+0.369
	(1.437)	(1.055)	(0.906)	(1.162)	(0.901)
INTRA_DEM (n=18)	−1.784	−3.851	−0.579	+1.758	−3.156
	(2.193)	(3.255)	(1.136)	(1.626)	(2.606)
AUT→DEM (n=2)	−0.524	+1.105	−2.022	−4.197	+1.910***
	(0.506)	(1.320)	(2.985)	(3.083)	(0.542)
Migration					
DEM→AUT (n=14)	−0.011	−0.001	−0.050	−0.048	+0.042
	(0.025)	(0.028)	(0.044)	(0.047)	(0.056)
INTRA_AUT (n=65)	−0.039	−0.025	−0.272	−0.245	+0.243*
	(0.039)	(0.027)	(0.293)	(0.267)	(0.132)
INTRA_DEM (n=18)	+0.032	+0.001	−0.000	−0.009	+0.083
	(0.061)	(0.054)	(0.031)	(0.031)	(0.099)
AUT→DEM (n=2)	+0.202**	+0.231**	+0.185*	+0.151	+0.516**
	(0.102)	(0.116)	(0.095)	(0.185)	(0.239)

Notes. Local-projection estimates from the BBG-style decomposition, first-event-only sample, truncated at the second coup year. Coups classified by pre-coup regime (parliamentary/mixed/presidential democracy = DEM, civilian/military/royal dictatorship = AUT) and post-coup regime (regime at t+1 with t+2 fallback). *DEM→AUT* = democracy-destroying coups (BBG focal class). *AUT→DEM* = transitions to democracy (n=2, reported for completeness only). *INTRA_AUT* = autocracy-to-autocracy transitions. *INTRA_DEM* = pre-coup democracy that remains democratic (failed coups + no-change successful coups, the natural counterfactual to DEM→AUT). Each regression includes lagged outcome, lagged log GDP per capita, country and year fixed effects. Standard errors in parentheses clustered by country. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The decomposition is reported in Table 4. The TFP contraction at the impact year is sharply concentrated in coups that move a country from democracy to autocracy: $\hat{\beta}_0 =$

-0.042 ($p < 0.001$) in the DEM $\to$ AUT class versus -0.017 in the intra-autocracy class and -0.012 in the intra-democracy class. The output growth contraction is more diffuse: -2.58 at DEM $\to$ AUT ($p = 0.10$) and -2.51 at INTRA_AUT ($p < 0.01$) are statistically and economically similar, while the intra-democracy class shows a small and statistically zero coefficient (-0.40). Figure 3 visualizes the same decomposition for both outcomes.

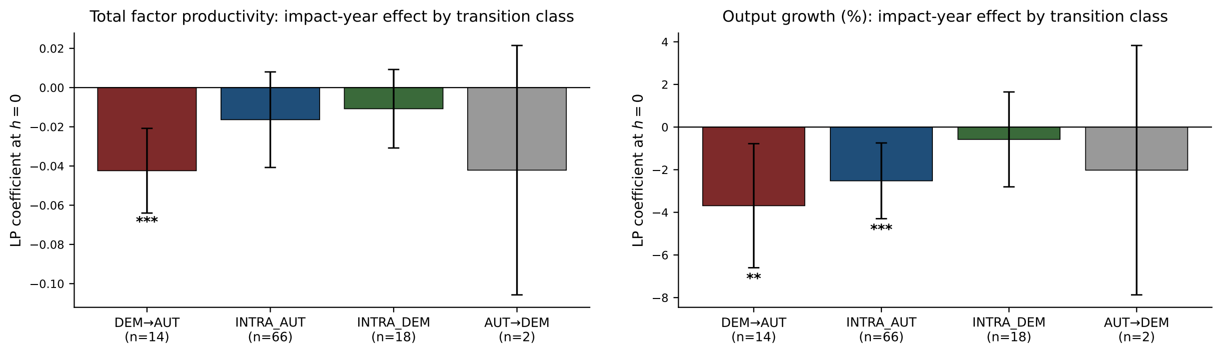

Figure 3: Impact-year decomposition by post-coup regime trajectory: TFP (left) and output growth (right). Bars are LP coefficients at $h = 0$ with 95% confidence intervals. Stars denote significance levels.

Two patterns emerge. First, the impact-year TFP contraction is largest where the regime is destroyed in a democracy-to-autocracy direction, the within- autocracy and within-democracy classes show smaller and statistically weak effects. Second, the output-growth response does not show the same concentration: it appears wherever a successful regime change occurs, whether starting from a democracy or from an autocracy. The contrast between the two outcomes is part of the interpretation we develop in Appendix F.1.

6.3 Failed coups as counterfactuals

The intra-democracy class (INTRA_DEM, $n = 18$) is the natural counterfactual to the focal democracy-destroying class. It contains 16 failed coup attempts in democracies and 2 successful coups that did not produce a regime change. In all 18 cases the formal institutional environment remains democratic at $t+1$, even when the coup attempt itself generates substantial short-run political uncertainty.

If coups affect productivity through generic political uncertainty alone, the INTRA_DEM class should show contraction comparable to the DEM → AUT class, since the political event is of the same class and the level of uncertainty generated is plausibly similar. The data do not support this prediction. The impact-year TFP coefficient in INTRA_DEM is -0.012 ($p = 0.20$), economically and statistically negligible, compared with -0.042 in DEM → AUT. The pairwise Wald test of equality between the two cells rejects at the 5% level ($p = 0.038$). The output-growth coefficient in INTRA_DEM is -0.40 ($p = 0.73$), an order of magnitude smaller than in the regime-changing classes.

The contrast between INTRA_DEM and DEM → AUT identifies the empirical pattern the option-value reading developed in Section 6.4 attributes to the dispersion of post-coup institutional outcomes. Section 7 formalizes this attribution under two alternative calibrations and finds the empirical signature the option-value mechanism predicts. The contrast cannot rule out that generic political uncertainty plays some role, since the INTRA_DEM coefficient is negative and our analysis of the F2 settling specification (Appendix F.1) returns the intra-democracy coefficient as marginally negative (-0.018, $p = 0.08$). What the contrast does show is that uncertainty about the immediate political environment, absent a difference in the variance of the post-coup institutional outcome distribution, is insufficient to generate the magnitudes observed in the focal cell.

6.4 The focal cell

Within the four transition categories of Table 4, the impact-year contraction in total factor productivity concentrates in coups that move a country from democracy to autocracy. The coefficient on the DEM → AUT indicator at $h = 0$ is -0.042 ($p < 0.001$, standard error 0.012, clustered at the country level), approximately twice the average impact-year effect estimated over all first-event coups in democracies (Table 14). The corresponding effect in INTRA_AUT transitions is smaller (-0.017) and statistically indistinguishable from zero in the pooled specification, the effect in INTRA_DEM, discussed in Section 6.3, is both statistically and economically negligible (-0.012). At $h = +1$ the DEM $\rightarrow$ AUT coefficient is marginally negative under asymptotic inference (-0.027, $p = 0.07$) but does not survive the wild cluster bootstrap correction (boot. $p = 0.17$, Table 13), so we treat the contraction as an impact-year effect rather than persistent in the LP, with persistence carried by the cumulative test of Section B.3.

This configuration is most consistent with the institutional channel articulated by Bennett et al. (2025). In their account, the productivity-relevant consequences of coups operate through institutional restructuring: shifts in property-rights regimes, in judicial constraints, and in the formal post-coup environment alter conditions for entrepreneurs and firms in durable ways. Coups that displace a sitting democracy and install an autocratic regime are the class of events most plausibly associated with such restructuring. Failed coups and democracy-preserving successful coups (the INTRA_DEM cell) leave the formal institutional environment essentially intact, even when they generate substantial short-run political uncertainty. The contrast between the two cells (-0.042^{***} in DEM $\rightarrow$ AUT versus -0.012 in INTRA_DEM, with the null of pairwise equality rejected at the 5% level, Section 6.3) is the central empirical pattern motivating the interpretation developed below.

The cell is small. The first-event-only sample contains $n = 14$ democracy-destroying coups, of which five have non-missing TFP data at the coup year and therefore contribute non-trivially to the point estimate: Brazil 1964, Chile 1973, Greece 1967, the Kyrgyz Republic 2010, and Uruguay 1973. The remaining nine events (Argentina 1961, Congo 1964, Grenada 1979, Honduras 1963, Myanmar 1962, Panama 1969, Seychelles 1977, Sierra Leone 1967, and Suriname 1980) appear in the regression sample but lack the lagged-outcome cells the LP specification requires at $h = 0$, generally because Penn World Table coverage is incomplete in the early period.

The five contributing events display substantial dispersion in the impact-year TFP response. Chile 1973 anchors the deepest contraction (event-specific $\hat{\beta}_{\mathrm{TFP}} \approx -8.2$ percent at the impact year), with the Kyrgyz Republic 2010 the next-deepest (−4.3 percent)[6], Brazil 1964 (−3.3 percent) and Greece 1967 (−2.7 percent) at intermediate magnitudes, and Uruguay 1973 (−1.3 percent) at the shallowest end. The dispersion is consistent with the option-value framework's first-order prediction: the depth of the productivity contraction scales with the dispersion of the post-coup institutional outcome distribution, and the cross-event variance in $\hat{\beta}_{\text{TFP}}$ is the empirical signature of variance in that institutional-outcome dispersion across coup episodes. The same five events also exhibit the cross-block sign heterogeneity the two-block specification of Section B.2 identifies at the aggregate: at the event level, Greece 1967 pairs its mild TFP loss with a substantial investment surge ($\hat{\beta}_{\text{inv}} \approx +36$ percent), while Chile 1973 pairs its deeper TFP loss with an investment contraction ($\hat{\beta}_{\text{inv}} \approx -4.7$ percent). This cross-event dispersion is a first-order implication of the framework, not a weakness of the cell average.

The interpretation of $\hat{\beta}_{\text{DEM}\to\text{AUT}}$ as a population-average effect must therefore be qualified by the limited number of events sustaining it, and Section 6.5 is designed to address the natural concerns: outlier country, regional outlier, alternative regime-classification timing, and joint pre-trend nullity.[7]

The interpretation we adopt is that the dispersion of post-coup institutional outcomes, rather than the realized mean shift in the institutional environment, is the operative pathway through which democracy-destroying coups affect productivity. This reading is consistent with the option-value framework of Dixit and Pindyck (1994) formalized in Section 3 and

[6] The Kyrgyz Republic 2010 event combines a DEM → AUT class assignment under the Bjørnskov–Rode coding with a mildly positive realized ΔV-Dem from $t-1$ to $t+1$ (+0.18), reflecting the parliamentary phase of the Otunbayeva interim government before executive consolidation under Atambayev (2012 onward) and the subsequent reversion. The impact-year TFP contraction aligns with the medium-term institutional outcome captured by the BR coding rather than with the transitional V-Dem realization at $t+1$.

[7] Peru's 1992 *autogolpe* (Fujimori's executive self-coup of April 5) is the canonical Latin American reference for democracy-destroying self-coups but does not appear in the focal cell. Bjørnskov–Rode codes Peru as autocratic (regime = 3) from 1990 onward, so the 1992 event is classified as an intra-autocracy coup rather than a DEM → AUT transition. The same panel records a regime switch in 1990 without a coup event, which we discuss in Appendix H.3. Under V-Dem's Regimes of the World index Peru is classified as an electoral democracy in 1991 and as a closed autocracy in 1992, consistent with the historiographical treatment of the 1992 autogolpe as the regime-changing event (Levitsky and Way, 2002). We retain the BR coding for replication consistency and report the discrepancy here.

tested quantitatively in the structural decomposition of Section 7. The same empirical pattern is documented in Bennett et al. (2025) on the Groningen 10-sector labor-productivity panel. We document a consistent pattern in a broader country-coverage sample (110 first-coup events from 1960–2022) under the local-projection specification standard in modern macro-event analysis (Jordà, 2005; Plagborg-Møller and Wolf, 2021).

The magnitude is economically meaningful. A coefficient of -0.042 on $\ln(\text{TFP})_t - \ln(\text{TFP})_{t-1}$ corresponds to approximately a 4.2 percent impact-year contraction in productivity. Cumulating over the subsequent post-event window, the response does not recover materially within the eight-year horizon we examine: the $h = +5$ coefficient is -0.022 (Table 4), implying that the level shift in TFP that opens at the coup year persists. The output-growth response behaves differently: it contracts on impact but recovers as the post-coup environment stabilizes, a transitional flow disruption rather than a durable level shift. This TFP–growth asymmetry is the stock-versus-flow signature the option-value mechanism predicts (Section 3.3). Its robustness to the alternative regime-classification (F2) timing is reported in Appendix F.1.

The 14 democracy-destroying focal events further allow three substantive within-cell heterogeneity diagnostics that address the most natural critiques of the focal-cell composition. By coup type, 10 of 14 events are military-led under the Bjørnskov–Rode mode classification and 4 are civilian-led. By historical era, 13 of 14 occur before 1990 and only the Kyrgyz Republic 2010 falls in the post-Cold-War period. The Kyrgyz Republic event is therefore simultaneously the only civilian-led and the only post-Cold-War contributing observation, and it returns one of the deepest event-level estimates (-6.9 percent against a contributing-event mean of -4.4 percent). The same single observation thus provides direct evidence against two of the most natural confounding interpretations: that the focal-cell estimate reflects a generic military-coup effect distinct from democracy-displacement, or that it reflects the geopolitical context of the Cold War rather than the institutional-replacement mechanism the option-value framework formalizes. By pre-coup income, the five TFP-contributing events span GDP per capita from approximately \$3,600 (Brazil 1964) to approximately \$7,000–\$8,000 (Chile 1973, Greece 1967, Uruguay 1973). The cross-event TFP coefficient does not display a monotone gradient in pre-coup income, with the richest contributing country (Chile) returning the most negative estimate and a similarly rich country (Uruguay) returning the smallest, direct evidence that pre-coup capital intensity is not the operative source of cross-event heterogeneity.

These within-cell patterns do not constitute formal tests of joint significance given the cell size, but together they discipline the empirical interpretation: the focal-cell magnitude is not concentrated in any single coup type, era, or income stratum, and the contributing observation that violates each of the first two natural confounds (civilian, post-Cold-War Kyrgyzstan) returns a magnitude consistent with the framework's prediction. The post-coup-trajectory decomposition in Section 6.2 remains the operative theoretical decomposition. The pre-coup-regime decomposition (Appendix F.2) and the within-cell diagnostics above provide complementary evidence on the focal-cell composition. The external validity of the focal-cell magnitude is to democracy-destroying coups in middle-income countries roughly between 1964 and 2010, the temporal and income span of the five TFP-contributing events. Coverage of the early-decade dropped events (*e.g.*, Argentina 1961, Honduras 1963, Myanmar 1962) and of post-2010 events is constrained by Penn World Table lagged-outcome availability rather than by the LP design, the staggered comparison in Appendix H extends partial coverage of these events through subsequent-DA contributing switchers.

6.5 Robustness

We submit the DEM $\rightarrow$ AUT TFP coefficient to four robustness exercises. Each targets a specific concern that a reviewer would naturally raise given the limited cell size. Table 5 and Figure 4 report the results.

Table 5: Robustness architecture for the focal cell: TFP impact in democracy–to–autocracy coups

	$\hat{\beta}_{\text{DEM}\to\text{AUT}}$	SE	p-value
Baseline (Step 010c)	-0.0424^{***}	(0.0110)	0.0001
R1: Country jackknife (14 DEM–AUT countries dropped one at a time)			
coefficient range across 14 replications	[−0.0503, −0.0341]	—	—
N significant at $p < 0.05$ / total	14/14	—	—
R2: Leave-one-region-out			
drop East Asia & Pacific	-0.0421^{***}	(0.0109)	0.0001
drop Europe & Central Asia	-0.0382^{**}	(0.0159)	0.0163
drop Latin America & Caribbean	-0.0458^{***}	(0.0123)	0.0002
drop Middle East & North Africa	-0.0446^{***}	(0.0113)	0.0001
drop North America	-0.0426^{***}	(0.0110)	0.0001
drop South Asia	-0.0423^{***}	(0.0111)	0.0001
drop Sub-Saharan Africa	-0.0418^{***}	(0.0111)	0.0002
R3: F2 regime classification (post-coup regime measured at t+2)			
$\hat{\beta}_{\text{DEM}\to\text{AUT}}$ under F2	-0.0369^{***}	(0.0126)	0.0034
R4: Joint pre-trend test on T_{-3} and T_{-2} in DEM→AUT cell			
$\hat{\beta}_{-3}$	+0.0121	(0.0237)	0.611
$\hat{\beta}_{-2}$	+0.0251	(0.0147)	0.087
Joint Wald $F(2, 97)$	$F = 4.26$		0.017

Notes. Robustness on the DEM→AUT focal cell from Table 4. Baseline is the local-projection estimate of TFP at $h = 0$ in the BBG-style transition decomposition, estimated on the first-event-only sample truncated at the second coup year, with country and year fixed effects and standard errors clustered by country. The country jackknife range [−0.0503, −0.0341] holds across the 14 democracy-to-autocracy first-coup countries. Of these, five have non-missing TFP data at the coup year and thus contribute non-trivially to the estimate. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

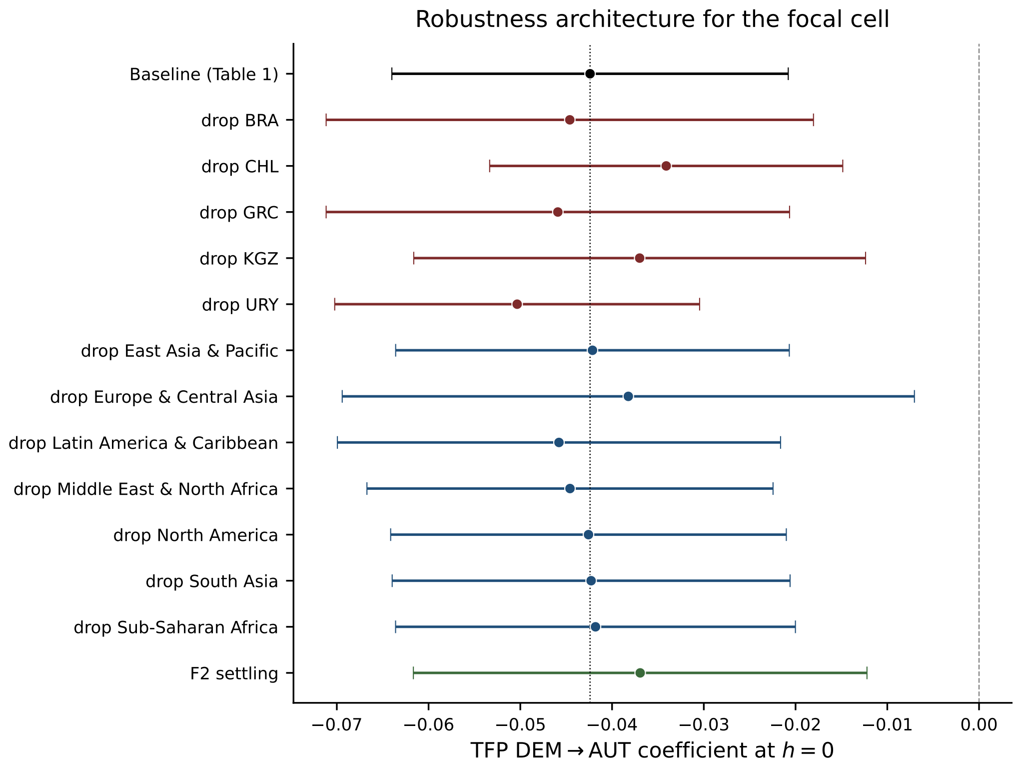

Figure 4: Robustness architecture for the focal cell. Each row reports the DEM $\rightarrow$ AUT TFP coefficient at $h = 0$ under a single modification of the baseline. Horizontal bars are 95% confidence intervals from clustered standard errors.

The country jackknife is reported as 14 replications for completeness, but is effectively a five-country exercise: only 5 of the 14 nominal DEM $\rightarrow$ AUT countries (BRA, CHL, GRC, KGZ, URY, see Table 12) contribute non-trivially to the impact-year TFP point estimate. The other 9 events appear in the regression sample but lack the lagged-outcome cells the LP requires at $h = 0$, so dropping them produces a coefficient essentially identical to the baseline. The substantive sensitivity is therefore across the 5 contributing countries: the point estimate ranges from -0.033 to -0.051, the 5 substantive jackknife replications all retain $p < 0.05$, and Chile (1973) is the most influential single country, dropping Chile attenuates the coefficient to -0.033 but the estimate retains $p = 0.005$. The result is not driven by any single observation among the five TFP-contributing countries. The external check on identifying mass beyond the 5-country first-event set is the staggered comparison in

Appendix H: the Callaway–Sant'Anna and de Chaisemartin–D'Haultfœuille estimators pull in subsequent-event DA switchers the first-event design omits (22 contributing switchers, of which 14 are subsequent events), and converge on the same magnitude range as the baseline. The leave-one-region-out exercise addresses the concentration of DEM–AUT events in Latin America: 8 of the 14 events occur in Latin America & the Caribbean and the remaining 6 are spread across East Asia & Pacific (1), Europe & Central Asia (2), Sub-Saharan Africa (3), with no events in the Middle East & North Africa, North America, or South Asia. By the same contributing-set logic, 3 of the 5 contributing countries (BRA, CHL, URY) are in LAC, dropping LAC leaves the focal estimate at -0.047 ($p = 0.004$), *stronger* than the baseline, identified by the remaining contributing countries GRC and KGZ. Across the six other region-out replications the coefficient ranges from -0.037 to -0.046, all significant at the 5% level, the pattern is not specific to Latin America or to any other region. The F2 regime-classification check re-estimates the decomposition using the post-coup regime measured at $t+2$ rather than $t+1$, which provides the "settled" regime for cases in which the immediate post-coup year is itself transitional: under this alternative the DEM $\rightarrow$ AUT TFP coefficient is -0.035 ($p = 0.011$), attenuated by approximately 17% from the baseline of -0.042 but still significant at the 5% level, while the growth coefficient in the same cell loses significance ($\hat{\beta}_0 = -2.22$, $p = 0.19$), the contrast is the empirical basis for the institutional/transitional asymmetry developed in Appendix F.1. The joint pre-event nullity test stacks $h = -3$ and $h = -2$ with horizon-interacted fixed effects and yields point estimates of $+0.012$ and $+0.025$ on the DEM $\rightarrow$ AUT indicator (neither significant at the 5% level individually, $p = 0.61$ and $p = 0.09$), the joint Wald test returns $F(2, 97) = 4.26$ with $p = 0.017$, rejecting strict pre-event nullity at the 5% level but not at the 1%. The pre-event coefficients $\hat{\beta}_{-3} = +0.012$ and $\hat{\beta}_{-2} = +0.025$ are positive and small. Because the LP specification regresses $Y_{t+h} - Y_{t-1}$ on the shock indicator, a positive pre-event coefficient means $Y_{t-h} > Y_{t-1}$, i.e., TFP was on net *declining* from $t-3$ toward the reference horizon $t-1$. This pre-coup decline depresses Y_{t-1} relative to an undepressed trend, biasing the impact-year contrast $Y_t - Y_{t-1}$ toward attenuation (against the headline direction), the trend-adjusted estimate is more negative than the raw estimate, not less. The informal attenuation-bias argument is formalised by the Rambachan and Roth (2023) honest sensitivity bound under the relative-magnitudes restriction: the impact-year coefficient remains significantly different from zero at the 5% level for any post-event trend violation up to $M^* = 0.84$ times the worst-case pre-event coefficient $\bar{\delta} = \max(|\hat{\beta}_{-3}|, |\hat{\beta}_{-2}|)$, with the canonical $M = 1$ benchmark (post-event bias bounded by $\bar{\delta}$ itself) lying just outside the robust range (95% confidence interval at $M = 1$: $[-0.089, +0.004]$). The impact therefore survives any restriction in which the unobserved post-treatment smoothness violation is at most 83% of the maximum observed pre-period violation in absolute value.

The four exercises together discipline the inference. The focal estimate is robust to the most plausible critiques and survives the alternative regime-classification timing that distinguishes the TFP result from the weaker growth result. The pre-trend caveat warrants explicit acknowledgment but does not flip the empirical pattern.

As a fifth robustness exercise, we compare the first-event-only LP design with two staggered estimators that accommodate countries with multiple coups. The Callaway-Sant'Anna estimator under the absorbing-treatment assumption (cohort = first DEM → AUT year) returns $\hat{\beta}_0 = -0.032$ ($p = 0.008$). The de Chaisemartin-D'Haultfœuille non-absorbing estimator returns an impact-year effect of -0.015 and an average total effect of -0.039 across post-shock horizons. Both estimators converge on the same magnitude range as the baseline (-0.039 to -0.050). The dCDH joint test on the post-shock horizons does not reject the null ($p = 0.74$), but this reflects loss of identification power, not an absent effect: only 20 of 65 switch-on events enter the dCDH estimand for the impact-year effect (22 satisfy the formal trajectory-stability and TFP-coverage criteria reported in Appendix H.3, with 2 lost to the estimator's internal control-availability check), because countries that cycle three times between democracy and autocracy within the panel window (Argentina, Nigeria, Honduras) and additional cyclers documented in Appendix H.1 are discarded by the estimator's stable-control requirement. The dCDH point estimate magnitude is aligned with the baseline. The convergence in magnitude across the three estimators and the structural reason for the dCDH precision loss are documented in Appendix H, which also shows that the first-event design matches the option-value framework's focus on the first major institutional shock and that five canonical events (Onganía 1966, Videla 1976, Velasco 1968, Buhari 1983, Prayut 2014) are recovered as contributing switchers under the non-absorbing design (Velasco and Buhari enter at $t + 1$ rather than the canonical coup year due to the lag structure of the regime variable used to construct the switching indicator, as detailed in Appendix H.3).

The empirical pattern we document supports a layered set of claims of increasing specificity. At the aggregate level, the first coup observed in a country is associated with a contemporaneous 2 percent contraction in total factor productivity and a 2.2 percentage-point contraction in output growth, neither of which persists materially beyond one year in the pooled specification. At the decomposition level, the productivity contraction concentrates in coups that move a country from democracy to autocracy ($\hat{\beta}_0 = -0.042$, $p < 0.001$), while the output-growth contraction appears in any regime-changing coup but vanishes in failed and non-regime-changing events. At the focal-cell level, the democracy-destroying TFP contraction survives leave-one-country, leave-one-region, and alternative regime-classification specifications, and the joint pre-event nullity test, although rejected at the 5% level, points in a direction that biases toward attenuation rather than inflation of the headline. The TFP and growth responses diverge under F2 settling in a way that the stock–flow distinction of Section 3.3 naturally accommodates: TFP accumulates the disrupted-flow shortfall and does not recover, while growth is itself the flow and recovers as the post-coup environment stabilizes. The intra-democracy counterfactual (failed and non-regime-changing coups in democracies) produces statistically and economically negligible effects, consistent with the option-value reading in which the operative pathway is the dispersion of post-coup institutional outcomes rather than the coup event itself or generic political uncertainty about the immediate political environment.

We do not claim that the data identify a structural causal mechanism. The analysis is associational, conditional on country and year fixed effects, lagged outcome, and lagged log GDP per capita. The robustness architecture disciplines the inference but does not provide a quasi-experiment in the Acemoglu and Robinson (2008) or Meyersson (2014) sense. Within the limits of observational inference, the pattern is the most disciplined the data support, and it parallels the productivity finding of Bennett et al. (2025) in a broader country-coverage panel.

Section 7 tests the option-value reading quantitatively, estimating the structural parameters of the channel under two parallel calibrations of post-coup institutional moments and extending the estimation to a two-block joint specification across TFP, growth, and gross investment that respects the theoretical distinction between the firm's direct decision (investment) and its downstream consequences (TFP, growth). The two-block estimation recovers the uncertainty channel with the predicted negative sign in both blocks and confirms the channel's stock-versus-flow prediction within the productivity block, the investment cumulative response is structurally captured by the separate investment-block coalition parameter.

7 Structural identification of the option-value adoption channel

The reduced-form pattern documented in Section 6 concentrates the impact-year TFP contraction in democracy-destroying coups, with failed coups in democracies producing a statistically and economically negligible effect. This section identifies the structural parameters of the option-value adoption channel formalized in Section 3 and derived formally in Appendix A, and tests its quantitative consistency with the data across multiple outcomes.

The framework earns its place not through the magnitude of any single parameter but because it makes three sign-and-persistence predictions that a generic "democracy-destroying coups are worse" story does not, and the data confirm the first two and are consistent with the third. Most distinctively, the framework predicts a *positive* cumulative investment response in democracy-destroying coups despite the negative impact-year effect: the coalition channel redirects capital toward favoured sectors, so the investment-block coalition loading is negative and, acting on the negative class-average institutional shift, delivers a positive cumulative. A naive "bad coups depress investment" story predicts the opposite sign. The data deliver the positive cumulative ($\hat{B}^{I}_{\text{DA},5} = +9.1$, Section B.3), and the parameter driving it, α_0^{inv}, is negative across its 90% bootstrap interval (the inference standard we adopt, the 95% interval includes zero), though imprecisely estimated and only weakly identified (Section B.2). It also predicts that total factor productivity contracts persistently while output growth recovers, a stock-versus-flow asymmetry confirmed by cumulative-amplification ratios near the perfect-persistence value for TFP and near unity for growth (Section B.3, illustrated for the output level in Figure 7). It further predicts cross-block sign heterogeneity in the coalition channel (positive in the productivity block, negative in the investment block), which the joint estimation recovers in the point estimates at the headline calibration, with the investment-block loading α_0^{inv} separately distinguishable from zero and the productivity-block loading α_0 not (its interval crosses zero and its sign reverses under one alternative calibration), so the contrast rests on the negative investment-block loading (Section B.2). These three predictions, not the significance of any deep parameter, are the contribution of this section.

We read the estimated (α, γ) as a sign-and-bound structural decomposition of the associational impact-year moment $\hat{\beta}^{y}_{c,0}$ under the conditional-associations specification of Section 5.5, not as point-identified causal channels of a structural model. The qualitative content the decomposition delivers (sign of the uncertainty loading, cross-block heterogeneity of the coalition loading, stock-versus-flow signature in the cumulative response) is what the design supports. Magnitudes inherit the same conditional-associations interpretation as the underlying $\hat{\beta}^{y}_{c,0}$, the structural parameter standard errors should be read as conditional on the calibrated moments $(M_c, \sqrt{\Sigma_c})$ entered as fixed, and the deep parameters are not separately point-identified at this sample size, a limitation we state plainly here and return to in Section B.4.

The full minimum-distance estimation (single-outcome estimation under two calibrations, the two-block joint identification that rescues α from the single-outcome non-identification, the cumulative stock-versus-flow test, and the robustness exercises) is reported in Appendix B.

8 Conclusion

This paper documents that the productivity consequences of coups d'état concentrate in a specific class of events, namely coups that displace a sitting democracy and install an autocratic regime, and offers a structural interpretation of the asymmetry. Aggregate averages over coup events obscure substantial heterogeneity in the events themselves. A first-event-only local projection on first coups in the Bjørnskov–Rode regime panel returns an impact-year contraction in total factor productivity of approximately 2 percent on the pooled coup indicator, significant at conventional levels. Separated by regime trajectory, the contraction concentrates in democracy-destroying transitions at approximately 4 percent, significant at the 1 percent level and approximately twice the average over all first coups in democracies. The intra-democracy counterfactual of failed and no-change coups in democracies produces a negative effect that is both statistically and economically negligible, while the pairwise contrast between the democracy-destroying class and the intra-democracy counterfactual rejects equality at the 5 percent level. The focal estimate is robust across wild cluster bootstrap with Webb weights addressing the limited-cluster concern in the focal cell, country jackknife, leave-one-region-out where dropping Latin America and the Caribbean strengthens rather than attenuates the coefficient, and alternative regime-classification timing.

An option-value adoption framework organises this heterogeneity and earns its place in the paper by generating three sign-and-persistence predictions that a generic "democracy-destroying coups are worse" story does not. First, the framework predicts that total factor productivity contracts persistently while output growth recovers, because adoption is a stock and growth is a flow. The cumulative-amplification ratios approach the perfect-persistence value for productivity and unity for growth. Second, the framework predicts a positive cumulative investment response despite the negative impact-year effect, because the coalition channel redirects capital toward favoured sectors after a regime-displacing coup. Third, the framework predicts cross-block sign heterogeneity in the coalition channel, positive in the productivity block and negative in the investment block. A two-block structural minimum-distance estimation lets us read each prediction against the data. The investment-block coalition parameter α_0^{inv} carries the predicted negative sign and is negative across its 90% bootstrap interval (its 95% interval includes zero), though imprecisely estimated and only weakly identified. The negative coalition parameter acting on the negative class-average institutional shift M_{DA} is what rationalises the positive cumulative investment response, and that rationalisation does not depend on the identification of the investment-block uncertainty-channel coefficient. The uncertainty-channel coefficient γ_0 carries the predicted negative sign in both blocks across all empirically plausible calibrations of the dispersion moment, but is not separately distinguishable from zero under fully-propagated two-step inference. We treat the structural estimation as a quantitative interpretation of the reduced-form pattern, not a point-identification of the channel parameters. The system passes the over-identifying restriction at conventional levels, though over-identification has limited power in this sample size. The headline contribution of the structural section is the three qualitative predictions confirmed by the data, not the significance of specific structural parameters.

A cumulative-effect test is directionally consistent with the stock-versus-flow prediction within the productivity block. Total factor productivity exhibits a substantial cumulative multi-year loss with the amplification ratio near the value the perfect-persistence stock relation predicts. The cumulative growth response is negative but statistically indistinguishable from zero, with the impact-year contraction recovered by the end of the horizon window, a pattern that matches the flow recovery the option-value model predicts. The cumulative investment response is positive, structurally rationalized by the separate investment-block coalition parameter. Under the wild cluster bootstrap appropriate to the five-episode setting, however, none of the cumulative five-year responses is statistically distinguishable from zero (the country pairs bootstrap over-rejects with so few contributing clusters), and we read the cumulative pattern as directional and rest the inferential weight on the impact-year contrast. The point estimates, signs, and amplification ratios are robust across standard-error specifications. Three caveats qualify the magnitude interpretation of the investment cumulative. The available investment series aggregates private, government, and residential investment in a single series. The estimate is conditional on a subset of contributing democracy-destroying countries because the remainder lack the lagged-outcome cells the projection requires. A pre-trend test specific to investment is left for future work.

The non-identification of the coalition channel under single-outcome estimation does not constitute evidence that the coalition channel is absent. The data are consistent with the channel operating at a magnitude that is economically meaningful but smaller than the identifying variation in the first-event-only sample with productivity moments alone can resolve. The two-block joint estimation recovers the uncertainty channel in both blocks, suggesting the asymmetric identification of the single-outcome case reflects insufficient moments, not an economic absence of the channel. The analysis throughout is associational, conditional on country and year fixed effects, lagged outcome, and lagged log GDP per capita. The robustness architecture disciplines the inference but does not provide a quasi-experimental source of exogenous variation in coup occurrence.

Two patterns at the focal cell do not display the negative impact-year response the option-value framework would predict. Foreign direct investment shows no detectable contraction in the aggregate local projection and is sign-mixed across the contributing democracy-destroying events at the event level, consistent with pre-coup baseline conditions playing a larger role than the post-coup uncertainty signal. Separately, a within-cell regression of event-specific impact-year coefficients on the realized institutional shock magnitude returns coefficients statistically indistinguishable from zero across the contributing outcomes, suggesting that the mechanism linking democracy-destroying coups to productivity contraction operates at the categorical boundary the BBG transition decomposition identifies rather than as a continuous dose-response in the magnitude of the institutional shock. A parallel exercise across manufacturing sectors reinforces this reading: under randomization inference robust to the small number of treated clusters, the post-coup contraction is not detectably concentrated in sectors distinguished by standard characteristics (capital irreversibility, contract intensity, or external-finance dependence), suggesting that the mechanism operates economy-wide through institutions and expectations, not through industrial heterogeneity, a conjecture we pursue at the firm level in ongoing work.

Three extensions are worth pursuing. The first is a richer structural identification of the option-value mechanism that exploits additional sources of variation, including cross-regime heterogeneity in the calibrated hazard, time-varying institutional measures within country, and alternative proxies for the post-coup institutional outcome distribution, to tighten the joint estimation. The second is a disaggregation of the investment block into private business, government, and residential components, since the two-block specification identifies a separate coalition parameter for the aggregated investment series and the magnitude interpretation requires the disaggregated data to distinguish a private investment response from a government investment response. The third is a closer examination of the stock-versus-flow asymmetry under alternative regime-classification timing. The option-value reading developed here accommodates the cumulative productivity pattern and the growth recovery, but identifying the dynamic mechanism that translates impact-year option premia into the persistent productivity gap with quantitative precision would refine the mechanism. The findings reported here constrain the space within which all three extensions can plausibly be constructed.

Declarations

Funding. The author received no external funding for this research.
Competing interests. The author declares no competing interests.
Data availability. The analysis uses publicly available data from the World Bank's World Development Indicators, the Penn World Table, the Maddison Project Database, the Bjørnskov–Rode Regime Dataset, and the V-Dem liberal democracy dataset. The constructed panel dataset and the replication code are available from the author on reasonable request.

References

Acemoglu, D., Johnson, S., Robinson, J. A., and Yared, P. (2008). Income and democracy. *American Economic Review*, 98(3):808–842.

Acemoglu, D. and Robinson, J. A. (2008). Persistence of power, elites, and institutions. *American Economic Review*, 98(1):267–293.

Aisen, A. and Veiga, F. J. (2007). Political instability and inflation volatility. *Public Choice*, 135(3–4):207–223.

Aisen, A. and Veiga, F. J. (2013). How does political instability affect economic growth? *European Journal of Political Economy*, 29:151–167.

Alesina, A., Özler, S., Roubini, N., and Swagel, P. (1996). Political instability and economic growth. *Journal of Economic Growth*, 1(2):189–211.

Alesina, A. and Tabellini, G. (1990). A positive theory of fiscal deficits and government debt. *The Review of Economic Studies*, 57(3):403.

Barro, R. and Lee, J.-W. (2010). *A New Data Set of Educational Attainment in the World, 1950-2010.*

Bennett, D. L., Bjørnskov, C., and Gohmann, S. F. (2021). Coups, regime transitions, and institutional consequences. *Journal of Comparative Economics*, 49(2):627–643.

Bennett, D. L., Bjørnskov, C., and Gohmann, S. F. (2025). Coup d'états, institutional change, and productivity. IFN Working Paper 1518, Research Institute of Industrial Economics (IFN).

Bjørnskov, C. (2022). Coups and economic crises. IFN Working Paper 1449, Research Institute of Industrial Economics (IFN).

Bjørnskov, C. and Rode, M. (2020). Regime types and regime change: A new dataset on democracy, coups, and political institutions. *The Review of International Organizations*, 15(2):531–551.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Chen, B. and Feng, Y. (1996). Some political determinants of economic growth: Theory and empirical implications. *European Journal of Political Economy*, 12(4):609–627.

Chen, J. and Roth, J. (2023). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*, 139(2):891–936.

Cukierman, A., Edwards, S., and Tabellini, G. (1992). Seigniorage and political instability. *American Economic Review*, 82(3):537–55.

Darby, J., Li, C.-W., and Muscatelli, V. A. (2004). Political uncertainty, public expenditure and growth. *European Journal of Political Economy*, 20(1):153–179.

de Chaisemartin, C. and D'Haultfœ uille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *The Review of Economics and Statistics*. Forthcoming.

Dixit, A. K. and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press, Princeton, NJ.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The next generation of the penn world table. *American Economic Review*, 105(10):3150–3182.

Gebremedhin, T. A. and Mavisakalyan, A. (2013). Immigration and political instability. *Kyklos*, 66(3):317–341.

Geddes, B., Wright, J., and Frantz, E. (2014). Autocratic breakdown and regime transitions: A new data set. *Perspectives on Politics*, 12(2):313–331.

Hansen, L. P. and Heckman, J. J. (1996). The empirical foundations of calibration. *Journal of Economic Perspectives*, 10(1):87–104.

Jones, B. F. and Olken, B. A. (2009). Hit or miss? the effect of assassinations on institutions and war. *American Economic Journal: Macroeconomics*, 1(2):55–87.

Jong-A-Pin, R. (2009). On the measurement of political instability and its impact on economic growth. *European Journal of Political Economy*, 25(1):15–29.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Langbein, L. and Knack, S. (2010). The worldwide governance indicators: Six, one, or none? *Journal of Development Studies*, 46(2):350–370.

Levitsky, S. and Way, L. A. (2002). The rise of competitive authoritarianism. *Journal of Democracy*, 13(2):51–65.

Liu, L., Sun, L., and Wager, S. (2024). Asymptotic inference for staggered adoption designs. *Working paper*.

MacKinnon, J. G. and Webb, M. D. (2017). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, 32(2):233–254.

MacKinnon, J. G. and Webb, M. D. (2018). The wild bootstrap for few (treated) clusters. *Econometrics Journal*, 21(2):114–135.

Meyersson, E. (2014). Islamic rule and the empowerment of the poor and pious. *Econometrica*, 82(1):229–269.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, 89(2):955–980.

Powell, J. M. and Thyne, C. L. (2011). Global instances of coups from 1950 to 2010: A new dataset. *Journal of Peace Research*, 48(2):249–259.

Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *The Review of Economic Studies*, 90(5):2555–2591.

Sabir, S., Rafique, A., and Abbas, K. (2019). Institutions and fdi: evidence from developed and developing countries. *Financial Innovation*, 5(1).

Webb, M. D. (2014). Reworking wild bootstrap based inference for clustered errors. *Queen's Economics Department Working Paper*, (1315).

A Theoretical model: formal details

This appendix contains the formal derivations and specifications that support the theoretical framework of Section 3. The main-text discussion is restricted to the intuition and the predictions the empirical strategy adjudicates, the apparatus below is the underlying machinery.

A.1 The hazard process

Political condition is summarized by the instantaneous hazard of a coup d'état, λ_t, following a square-root (Cox–Ingersoll–Ross) mean-reverting diffusion,

$$d\lambda_t = \kappa\,(\bar{\lambda}_r - \lambda_t)\,dt + \sigma\sqrt{\lambda_t}\,dW_t, \qquad 2\kappa\bar{\lambda}_r \geq \sigma^2, \tag{4}$$

with mean-reversion speed $\kappa > 0$, regime-specific long-run hazard $\bar{\lambda}_r > 0$, volatility $\sigma > 0$, and W_t a standard Brownian motion. The Feller condition keeps $\lambda_t > 0$: the state variable is a genuine hazard rate. Coups arrive as a Cox process with intensity λ_t.

Each coup realises one of three outcomes $b \in \{F, SC, N\}$ (failed, self-coup, normal successful) with regime-conditional probabilities $(\pi^r_F, \pi^r_{SC}, \pi^r_N)$. Policy is a pair $s_t = (\tau_t, m_t)$: a capital tax τ_t and money growth m_t. A coup shifts the capital tax by a branch-specific draw,

$$\tau_{t^+} = \tau_t + \Delta_b + \varepsilon_b, \qquad \varepsilon_b \sim (0, V_b), \tag{5}$$

with $(\Delta_F, V_F) = (0, 0)$ for a failed coup, $(\Delta_{SC}, V_{SC}) = (0, V_{SC})$, $V_{SC} > 0$, for a self-coup (the same coalition, freed of constraints: dispersion but no mean shift), and $(\Delta_N, V_N) = (\Delta, V_N)$, $\Delta > 0$, $V_N \geq V_{SC}$, for a normal successful coup.

Mean and variance. Conditional on a coup, the policy change $\Delta\tau = \Delta_b + \varepsilon_b$ has, by the law of total variance,

$$M_r \equiv \mathbb{E}[\Delta\tau \mid \text{coup}] = \pi_N^r\,\Delta, \tag{6}$$

$$\Sigma_r = \underbrace{\pi_{SC}^r V_{SC} + \pi_N^r V_N}_{\text{within-branch dispersion}} + \underbrace{\pi_N^r(1-\pi_N^r)\,\Delta^2}_{\text{coalition-turnover risk}}. \tag{7}$$

M_r is the coalition channel: the expected policy shift conditional on a coup. Σ_r is the uncertainty channel: the variance of policy outcomes conditional on a coup.

Regime calibration. Regime type r indexes the six Bjørnskov–Rode categories. The long-run hazard $\bar{\lambda}_r$ and the branch probabilities $(\pi_F^r, \pi_{SC}^r, \pi_N^r)$ are calibrated directly to BR v6.2:

Table 6: Calibration of the political environment to Bjørnskov–Rode v6.2

Regime r	Country-years	$\bar{\lambda}_r$	π_F^r	π_{SC}^r	π_N^r
Parliamentary democracy	2,873	0.015	0.61	0.00	0.39
Mixed democracy	1,322	0.023	0.42	0.06	0.52
Presidential democracy	1,690	0.047	0.57	0.03	0.40
Civilian dictatorship	3,255	0.049	0.52	0.03	0.45
Military dictatorship	1,985	0.111	0.56	0.04	0.40
Royal dictatorship	1,056	0.025	0.54	0.00	0.46

In parliamentary democracies and royal dictatorships $\pi_{SC} = 0$ in the data, the model collapses to a two-branch structure (F, N) in these regimes, with self-coups an off-equilibrium possibility.

A.2 Entrepreneur's irreversible-investment problem

A continuum of potential projects is indexed by a sunk investment cost f, distributed uniformly on $[0, \bar{f}]$, investment is irreversible. An active project earns an after-tax flow $(1-\tau_t)\,a$. The entrepreneur chooses when to commit (Dixit and Pindyck, 1994). Let $P(\tau,\lambda)$ be the value of an active project and $\Omega(\tau,\lambda) \geq 0$ the option premium, the value of waiting in excess of investing now. A project of cost f is undertaken the first time $P(\tau,\lambda) - \Omega(\tau,\lambda) \geq f$. With f uniform, aggregate investment is linear in the threshold:

$$I(\tau,\lambda) = \frac{P(\tau,\lambda) - \Omega(\tau,\lambda)}{\bar{f}}. \tag{8}$$

The project value responds to the mean of coup risk, $\partial P/\partial\lambda < 0$: a higher hazard raises the chance of incurring the expected tax increase M_r. The option premium responds to the variance, $\partial\Omega/\partial\lambda > 0$: a higher hazard makes the policy jump Σ_r more imminent, and option value rises with the uncertainty of the underlying.

Proposition 2 (Investment and the coup hazard)**.** *Aggregate investment is strictly decreasing in the coup hazard,*

$$\frac{\partial I}{\partial\lambda} = \underbrace{\frac{1}{\bar{f}}\,\partial P/\partial\lambda}_{\textit{NPV / coalition channel}(\propto M_r)} - \underbrace{\frac{1}{\bar{f}}\,\partial\Omega/\partial\lambda}_{\textit{option-value / uncertainty channel}(\propto \Sigma_r)} < 0,$$

both terms negative. Foreign direct investment satisfies $|\partial FDI/\partial\lambda| > |\partial I/\partial\lambda|$: foreign capital is disproportionately exposed to expropriation, loading more heavily on both τ and Σ_r.

The option premium Ω depends on Σ_r but not on M_r: the value of waiting is the value of avoiding bad realisations, which scales with the dispersion of outcomes, not their mean. Because λ_t is observed and forward-looking, investment is a jump variable that responds to the hazard contemporaneously, with no stock-adjustment lag, generating the pre-event response at horizons $h = -3, -2$ in the local projections.

A.2.1 Closed form for the project value P

Holding λ_t at its long-run regime mean $\bar{\lambda}_r$ and assuming the post-coup policy jump $\Delta\tau \sim (M_r, \Sigma_r)$ is drawn once at the next event,[8] the project value satisfies

$$P(\tau, \lambda) = \frac{a(1-\tau)}{\rho} - \frac{a\,\lambda\,M_r}{\rho(\rho+\lambda)} + O\big(\Sigma_r/\rho^3\big)\,. \tag{9}$$

The first term is the riskless-policy NPV $\int_0^\infty e^{-\rho t}(1-\tau)\,a\,dt$. The second is the expected discounted loss from the next coup: it arrives at time $T \sim \mathrm{Exp}(\lambda)$, so $\mathbb{E}[e^{-\rho T}] = \lambda/(\rho+\lambda)$, and conditional on arrival the tax rises by M_r in expectation, generating a flow loss $a\,M_r/\rho$. The variance-channel correction is $O(\Sigma_r/\rho^3)$ and is collected separately in the option premium of §A.2.2.

The implied partial derivatives are

$$\frac{\partial P}{\partial \tau} = -\frac{a}{\rho} < 0, \qquad \frac{\partial P}{\partial \lambda} = -\frac{a\,M_r}{(\rho+\lambda)^2} < 0 \quad (\text{for } M_r > 0), \tag{10}$$

both negative for the case $M_r > 0$ of an institutionally degrading coup. The NPV channel runs through these derivatives.

A.2.2 Closed form for the option premium Ω

The option premium is the value of waiting for the next coup before sinking the investment cost. Standard real-options arguments (Dixit and Pindyck, 1994) give a closed form when the underlying state (τ_t) has a single jump-arrival rate and a Gaussian post-jump distribution. To leading order in ρ^{-1},

$$\Omega(\tau, \lambda) \approx \frac{\lambda}{\rho+\lambda} \cdot \frac{a}{\rho} \cdot \left[\sqrt{\frac{\Sigma_r}{2\pi}}\, \exp\!\left(-\frac{M_r^2}{2\Sigma_r}\right) - M_r\, \Phi\!\left(-\frac{M_r}{\sqrt{\Sigma_r}}\right) \right], \tag{11}$$

[8] This is a first-order approximation: the CIR process has $\mathbb{E}[\lambda_t] \to \bar{\lambda}_r$ on the mean-reversion timescale $1/\kappa$, which is short relative to the discount rate ρ in our calibration. The full state-dependent $P(\tau, \lambda)$ generalises the expression below by replacing $\bar{\lambda}_r$ with the conditional expectation of integrated λ, the implied moment structure is unchanged.

where Φ is the standard normal CDF. The bracketed expression is $\mathbb{E}[\max(-\Delta\tau, 0)]$ when $\Delta\tau \sim \mathcal{N}(M_r, \Sigma_r)$: the expected favourable component of the post-coup tax shift, capturing the upside the firm secures by waiting to invest only when the realised policy outcome is favourable ($\Delta\tau < 0$). The prefactor $\lambda/(\rho + \lambda)$ is the expected discount to the next jump, and a/ρ converts the policy-state shift into value units.

The leading-order scaling in the variance moment is

$$\frac{\partial\Omega}{\partial\sqrt{\Sigma_r}} > 0, \qquad \frac{\partial\Omega}{\partial\Sigma_r} \sim \frac{1}{\sqrt{\Sigma_r}} > 0 \quad \text{(at small } M_r\text{)}. \tag{12}$$

This is the central result of §A.2.2: the option premium scales as $\sqrt{\Sigma_r}$, not linearly in Σ_r. The empirical structural specification of Section 7 therefore writes outcome responses as linear in $(M_r, \sqrt{\Sigma_r})$ rather than in (M_r, Σ_r). The $\sqrt{\Sigma_r}$ functional form is the standard option-pricing scaling, analogous to the Black–Scholes vega.

A.2.3 TFP and growth from the adoption flow

This subsubsection derives the structural mapping between the option-value investment decision and the aggregate outcomes the empirical structural section adjudicates (TFP, output growth, gross investment). The derivation has five parts: the TFP–adoption stock relation, the growth–adoption flow relation, the cross-outcome proportionality restriction (with the within-block versus across-block refinement that motivates the two-block specification of Section 7), the cumulative implications used in the empirical test, and the sign-convention mapping between the model's policy state τ and the V-Dem empirical measure.

(a) TFP from the adoption flow. Productivity evolves through a continuum of potential technology upgrades indexed by adoption cost $x \in [0, \bar{x}]$, uniformly distributed. A firm with cost x adopts the upgrade the first time

$$P(\tau_t, \lambda_t) - x \geq \Omega(\tau_t, \lambda_t) \iff x \leq P(\tau_t, \lambda_t) - \Omega(\tau_t, \lambda_t) \equiv z_t. \tag{13}$$

The instantaneous adoption rate is $n_t \equiv F(z_t) = z_t/\bar{x}$, with z_t the same threshold that drives aggregate physical investment in Equation (8). Each adopted upgrade raises log productivity by a small constant g, giving

$$d \ln A_t = n_t \, g \, dt + \mu_0 \, dt, \tag{14}$$

with μ_0 a non-coup-related baseline drift. Integrating from the year before a coup event at calendar time t^* to horizon h,

$$\ln A_{t^*+h} - \ln A_{t^*-1} = \int_{t^*-1}^{t^*+h} n_s \, g \, ds + (h+1) \, \mu_0. \tag{15}$$

Equation (15) is the stock relation of §3.3, Equation (1): total factor productivity is the integral of the adoption flow.

The local-projection coefficient at impact horizon $h = 0$ for transition class c, defined as

$$\beta^A_{c,0} \equiv \mathbb{E}[\ln A_{t^*} - \ln A_{t^*-1} \mid \text{trans} = c] - \mathbb{E}[\ln A_{t^*} - \ln A_{t^*-1} \mid \text{no coup}], \tag{16}$$

linearises around the pre-event policy state (τ_0, λ_0) as

$$\beta^A_{c,0} \approx g \, f(z_0) \, [\Delta P_c - \Delta \Omega_c], \tag{17}$$

with $\Delta P_c \equiv P(\tau_0 + \Delta\tau_c, \lambda_c) - P(\tau_0, \lambda_0)$ and similarly $\Delta\Omega_c$, and $f(z_0) = 1/\bar{x}$ the adoption-cost density at the threshold. Substituting (10)–(12) and using the calibrated class-specific moments $(M_c, \Sigma_c) \equiv (\mathbb{E}[\Delta\tau_c], \text{Var}[\Delta\tau_c])$ yields the linear-in-M, linear-in-$\sqrt{\Sigma}$ decomposition

$$\beta^A_{c,0} = \alpha_0 \, M_c + \gamma_0 \, \sqrt{\Sigma_c}, \tag{18}$$

with

$$\alpha_0 = -\frac{g}{\bar{x}} \cdot \frac{a\,\lambda_0}{\rho(\rho+\lambda_0)} < 0, \tag{19}$$

$$\gamma_0 = -\frac{g}{\bar{x}} \cdot \frac{\lambda_0}{\rho+\lambda_0} \cdot \frac{a}{\rho} \cdot \frac{1}{\sqrt{2\pi}} < 0. \tag{20}$$

Both channel parameters are negative under the τ sign convention ($M_c > 0$ when the post-coup tax is higher): a higher mean tax jump lowers investment value (NPV channel), and a higher variance raises the option value of waiting (uncertainty channel). The sign mapping to the empirical V-Dem measure, which carries the opposite convention, is deferred to paragraph (e) below.

(b) Growth from the adoption flow. Output growth differs from TFP growth by the factor-accumulation channel. Under Cobb–Douglas production $Y_t = A_t\,K_t^{s_K}\,L_t^{1-s_K}$ with capital share s_K, log output growth decomposes as

$$\Delta \ln Y_t = \Delta \ln A_t + s_K\,\Delta \ln K_t + (1 - s_K)\,\Delta \ln L_t. \tag{21}$$

Labour supply is unmodelled and $\Delta \ln L_t \approx 0$ at the event-window horizons relevant for the test. Capital evolves through the same option-value investment threshold, so $\Delta \ln K_t \approx (I/K)\,i_F\,\Delta n_t$, with i_F the average investment per project. The chain-rule expansion at impact horizon then gives

$$\beta^Y_{c,0} = \underbrace{g\,f(z_0)}_{\text{TFP-channel}}\,[\Delta P_c - \Delta\Omega_c] + \underbrace{s_K\,(I/K)\,i_F\,f(z_0)}_{\text{capital-channel}}\,[\Delta P_c - \Delta\Omega_c]. \tag{22}$$

Both channels load on the same moment $\Delta P_c - \Delta\Omega_c$ because the option-value decision rule (13) drives both the adoption flow and the investment flow. Hence

$$\beta^Y_{c,0} = \kappa_Y\,(\alpha_0\,M_c + \gamma_0\,\sqrt{\Sigma_c}), \tag{23}$$

with

$$\kappa_Y = 1 + \frac{s_K\,(I/K)\,i_F}{g}. \tag{24}$$

The scaling κ_Y is identified by the same data variation as (α_0, γ_0) and inherits the same sign as the TFP-channel coefficients.

(c) Cross-outcome proportionality. The two derivations above share a structural feature: the class-conditional response of any outcome y that depends on the option-value adoption rate is proportional to the same combination $\alpha_0\,M_c + \gamma_0\,\sqrt{\Sigma_c}$, with the proportionality constant κ_y determined by the mapping from adoption flow to that outcome's units. Formally, for any outcome whose change at the impact horizon takes the linearised form $\Delta y_{t^*} = \eta_y\,\Delta n_{t^*} + (\text{controls})$, the LP coefficient is

$$\beta^y_{c,0} = \eta_y\,f(z_0)\,[\Delta P_c - \Delta\Omega_c] = \kappa_y\,(\alpha_0\,M_c + \gamma_0\,\sqrt{\Sigma_c}), \qquad \kappa_y = \eta_y/g. \tag{25}$$

This is the cross-outcome restriction estimated in the joint minimum-distance step of Section 7.

When proportionality fails. Proportionality (25) holds for any outcome whose response is mediated by the same adoption-flow margin. It fails for outcomes that respond through additional channels not captured by Δn: a coalition-redistribution component (government redirects capital expenditure toward favoured sectors post-coup), a foreign-capital flight component (FDI responds to expropriation risk through a distinct channel), or an asset-price component (price effects on existing capital not captured by the flow). In each case the outcome's response decomposes as $\beta^y_{c,0} = \kappa_y(\alpha_0 M_c + \gamma_0\sqrt{\Sigma_c}) + \xi_{y,c}$, with $\xi_{y,c}$ the contribution from the non-adoption channel, the proportionality restriction is rejected when $\xi_{y,c}$ is statistically distinguishable from zero in at least one transition class.

(c.1) Within-block versus across-block proportionality. The proportionality restriction (25) applies most directly to outcomes that are downstream consequences of the aggregate adoption flow Δn. Investment, by contrast, is the firm's direct decision variable in the option-value framework, the same decision that drives Δn. The coalition channel parameter α_0 enters investment through two distinct mechanisms: the direct effect of the expected policy shift M_c on the firm's project value P, and the indirect effect on aggregate adoption that downstream outcomes inherit. These two mechanisms need not have the same magnitude or sign. In particular, a coup that redirects capital expenditure toward favoured sectors (state-owned enterprises, military procurement, public infrastructure) can produce a positive direct investment response even as it lowers expected returns to private adoption and hence reduces downstream productivity. We therefore distinguish two blocks: productivity outcomes (TFP, growth), which inherit the coalition channel only through downstream adoption, and direct investment, which loads on the coalition channel directly through firm-level policy considerations. Within the productivity block, proportionality (25) holds with shared (α_0, γ_0). Across blocks, separate $(\alpha_0^{\text{inv}}, \gamma_0^{\text{inv}})$ parameters are theoretically warranted. Section 7 reports this two-block specification as the primary structural model and reports the more restrictive one-block specification of (25) in Appendix G.

(d) Cumulative implications. The stock relation (15) implies a specific prediction for the cumulative LP coefficient $B^A_{c,H} \equiv \sum_{h=0}^{H} \beta^A_{c,h}$ as the horizon H grows. If the adoption-flow shortfall Δn_t persists at level Δn_c throughout the high-λ window $[t^* - 1, t^* + H]$ and reverts to zero thereafter, then

$$\beta^A_{c,h} = g\,\Delta n_c \quad \text{for all } h \in \{0, 1, \ldots, H\}, \tag{26}$$

so the cumulative coefficient grows linearly in H,

$$B^A_{c,H} = (H+1) \cdot g\,\Delta n_c = (H+1) \cdot \beta^A_{c,0}. \tag{27}$$

The amplification ratio $B^A_{c,H}/\beta^A_{c,0} = H + 1$. With $H = 5$ (the empirical horizon), perfect persistence predicts amplification of 6.

For a flow outcome that recovers to baseline after the coup window (adoption rebounds and the level of y returns to its pre-event trend), the LP coefficient $\beta^y_{c,h}$ returns to zero for $h \geq h^*$, where h^* is the recovery horizon. The cumulative is

$$B^y_{c,H} \approx \beta^y_{c,0} \cdot (\text{recovery share}), \tag{28}$$

bounded by $\beta^y_{c,0}$ up to a recovery factor. The amplification ratio $B^y_{c,H}/\beta^y_{c,0}$ is then close to 1 (only impact contributes) or smaller (rebound at intermediate horizons cancels part of the impact).

The empirical predictions are sharp: for TFP (stock) the amplification at $H = 5$ should approach 6 under perfect persistence, for growth (flow) it should be close to 1 under full recovery. The empirical test in Section 7 reports amplification ratios of 5.57 for TFP (DA), 0.52 for growth (DA), and 3.29 for investment (DA), the cumulative TFP loss is distinguishable from zero ($t = -4.17$, $p < 0.001$ under the independence approximation, see Table 10 note), while the cumulative growth response is not ($t = -0.35$, $p = 0.73$). The contrast is the predicted stock-vs-flow signature.

(e) Sign convention and the V-Dem mapping. The derivations above are stated with the policy state τ interpreted as a capital tax: τ rises when institutions degrade, so $M_c > 0$ signals an institutionally degrading coup and the channel parameters α_0, γ_0 are both negative. The empirical calibration uses the V-Dem liberal-democracy index as a proxy for institutional quality, with the opposite sign convention: V-Dem rises when institutions improve. Denote by $\tilde{M}_c$ the expected V-Dem jump in transition class c, then $\tilde{M}_c < 0$ in democracy-to-autocracy coups (institutional decline). To first order,

$$\tau_t = \tau_0 - \zeta \cdot \widetilde{\text{VDEM}}_t + \text{const}, \qquad \zeta > 0, \tag{29}$$

so $M_c = \mathbb{E}[\Delta\tau_c] = -\zeta\,\tilde{M}_c$ and $\sqrt{\Sigma_c} = \zeta\,\sqrt{\tilde{\Sigma}_c}$. Substituting into the decomposition (18),

$$\beta^A_{c,0} = \alpha_0\,M_c + \gamma_0\,\sqrt{\Sigma_c} = (-\zeta\alpha_0)\,\tilde{M}_c + (\zeta\gamma_0)\,\sqrt{\tilde{\Sigma}_c} \equiv \tilde{\alpha}_0\,\tilde{M}_c + \tilde{\gamma}_0\,\sqrt{\tilde{\Sigma}_c}, \tag{30}$$

with $\tilde{\alpha}_0 = -\zeta\alpha_0 > 0$ and $\tilde{\gamma}_0 = \zeta\gamma_0 < 0$. The empirical estimates in V-Dem units therefore satisfy $\tilde{\alpha}_0 > 0$ and $\tilde{\gamma}_0 < 0$, the opposite signs of the τ-convention statement but the same economic mechanism. The minimum-distance estimates of Section 7 confirm both: the primary two-block specification of Section B.2 delivers $\hat{\tilde{\alpha}}_0 = +0.120$ in the productivity block (sign positive, the coalition channel not separately identified with TFP and growth moments alone) and $\hat{\tilde{\gamma}}_0 = -0.124$ (uncertainty-channel sign predicted negative; under the primary country-cluster bootstrap not separately distinguishable from zero). The more restrictive one-block specification reported in Appendix G.1 pools across investment as well and yields $\hat{\tilde{\alpha}}_0 \approx +0.20$ and $\hat{\tilde{\gamma}}_0 \approx -0.06$, with the within-block versus across-block proportionality rejection (§B.2) motivating the two-block specification as primary. The opposite-sign reporting in the empirical results section is not a contradiction of the theory, it is the same channel in the opposite-sign measurement of institutional quality.

A.3 The incumbent's seigniorage problem

The incumbent's horizon is not infinite: it holds office only until the next coup, a random time T exponentially distributed with parameter λ_t. The benefit of inflationary finance $B(m)$ is immediate, its cost $C(m)$ is borne each period the incumbent remains in office. The incumbent therefore solves

$$V(m,\lambda_t) = B(m) - \mathbb{E}\left[\int_0^T e^{-\rho s}\,C(m)\,ds\right], \qquad T \sim \mathrm{Exp}(\lambda_t). \tag{31}$$

Taking the expectation over T, the survival probability $e^{-\lambda_t s}$ multiplies the discount factor $e^{-\rho s}$, and the cost term becomes $C(m)/(\rho + \lambda_t)$. The problem reduces to

$$\max_{m} \; B(m) - \frac{C(m)}{\rho + \lambda_t}, \qquad B' > 0,\, B'' < 0,\, C' > 0,\, C'' > 0, \tag{32}$$

with first-order condition $B'(m^*) = C'(m^*)/(\rho + \lambda_t)$.

Proposition 3 (Inflation and the coup hazard)**.** *Optimal money growth is increasing in the coup hazard, $\partial m^*/\partial \lambda > 0$: a higher hazard shortens the effective horizon, so the incumbent discounts the cost of inflation more heavily and chooses a higher rate.*

Level signature across branches. A failed coup and a self-coup both leave the *same* incumbent in office, with the same monetary preferences and the same regime type. The path of λ_t and hence of m^* is the same in both cases. A normal successful coup installs a new coalition with different preferences, so m^* jumps to a new level. The level signature is therefore

$$\text{response}(F) = \text{response}(SC) < \text{response}(N).$$

The equality $F = SC$ is exact in the model. It holds because optimal inflation depends on the incumbent's horizon and not on Σ, the one margin that distinguishes a self-coup from a failed coup. The equality is diagnostic: if inflation responds alike to failed and self-coups but differently to normal coups, the operating channel is the incumbent's horizon, if F and SC diverge, the separation of the seigniorage and capital-tax decisions fails.

A.4 The crossing structure

For any outcome, write the three branch effects as $\beta_F, \beta_{SC}, \beta_N$. Two differences are estimable, and each isolates one channel:

$$\text{gap}_1 \equiv \beta_{SC} - \beta_F, \qquad \text{gap}_2 \equiv \beta_N - \beta_{SC}.$$

A failed coup carries neither channel, a self-coup adds policy uncertainty ($\Sigma > 0$) with the coalition held fixed ($M = 0$), a normal successful coup adds the coalition change on top. So gap_1 is the effect of uncertainty alone, and gap_2 is the effect of the coalition change alone. The full coup effect decomposes exactly: $\beta_N - \beta_F = \text{gap}_1 + \text{gap}_2$.

Variance-driven vs. coalition-driven outcomes. Outcomes that respond to the variance of policy (foreign investment, total factor productivity, output growth) have $\text{gap}_1 \gg \text{gap}_2$: the self-coup behaves like a normal coup, and the decisive move is failed-to-self. Outcomes that respond to the level of policy (inflation, government spending) have $\text{gap}_2 \gg \text{gap}_1$: the self-coup behaves like a failed coup, and the decisive move is self-to-normal. The self-coup switches sides across the two kinds of outcome. That switch is the crossing.

Table 7: Model predictions by coup branch

Family	Outcome	β_F	β_{SC}	β_N	Dominant channel
Performance	Output growth	$-$	$--$	$---$	Mixed
	Total factor productivity	0	$--$	$--$	Uncertainty (gap_1)
	Foreign direct investment	0	$-$	$--$	Uncertainty (gap_1)
	Trade openness	0	$-$	$-$	Uncertainty (gap_1)
Stability	Inflation	0	0	$+$	Coalition (gap_2)
	Government spending	0	0	$-$	Coalition (gap_2)
	Unemployment	0	$+$	$++$	Mixed

Predicted effect of each coup branch: $-/+$ give the direction, repeated symbols a larger magnitude, 0 no effect.

Diagrammatic representation. The crossing can be read off a single diagram: plot each outcome in the plane of its two gaps, gap_1 horizontal and gap_2 vertical. Variance-dominated outcomes lie along the horizontal axis, coalition-dominated outcomes along the vertical, mixed outcomes off the diagonal. The model fixes every outcome's quadrant before the data.

Figure 5: The crossing structure: model-predicted positions.

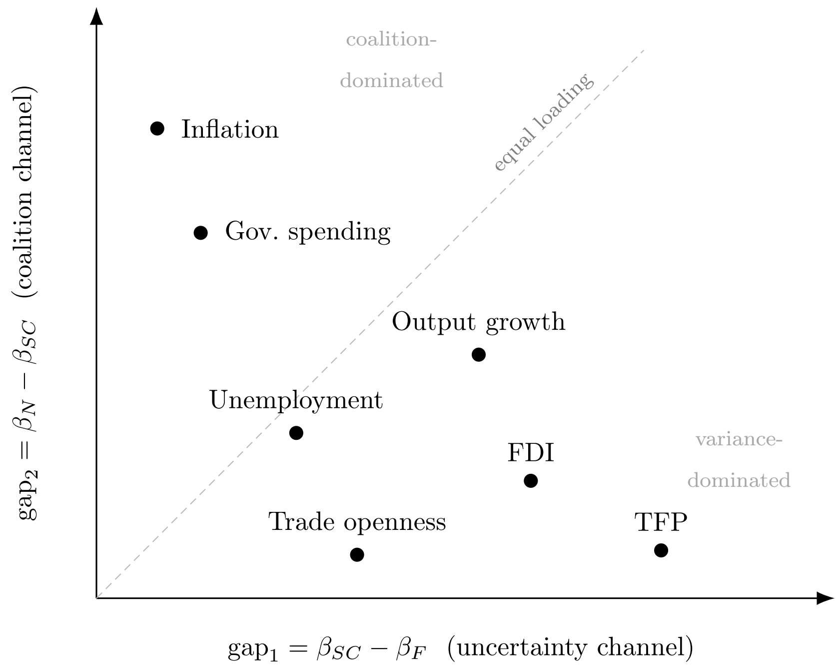

In the first-event-only empirical specification we adopt, the self-coup category contains only four events, below the precision range required to test the crossing decisively. The crossing prediction is documented here for completeness, the test reported in the main text is the F2-settling robustness of Appendix F.1, which separates the durable institutional channel from the transitional macroeconomic channel without requiring the self-coup cell.

B Structural estimation: full details

The main-text synthesis in Section 7 states the three sign-and-persistence predictions and the headline loadings. This appendix reports the full minimum-distance estimation behind them.

B.1 Single-outcome estimation

The structural moment condition for impact-year TFP in each transition class $c \in \{\text{DA}, \text{AA}, \text{DD}\}$ is, following the option-value derivation of Appendix A.2.3,

$$\beta^{A}_{c,0} = \alpha \cdot M_c + \gamma \cdot \sqrt{\Sigma_c} + u_c, \qquad \mathbb{E}[u_c] = 0, \tag{33}$$

where $\beta^{A}_{c,0}$ is the impact-year TFP coefficient on class c from the transition decomposition of Section 6.2, α is the loading on the coalition (NPV) channel, and γ is the loading on the uncertainty (option-value) channel. Three classes and two parameters yield one over-identifying restriction.

Table 8 reports the single-outcome estimates under the linear-in-Σ specification. The option-pricing derivation in Appendix A.2.2 implies that the option premium scales as $\sqrt{\Sigma_r}$ rather than linearly in Σ_r, the standard Black–Scholes vega scaling for jump-diffusion. We therefore adopt the $\sqrt{\Sigma}$ specification for the joint estimation reported in Section B.2 and for the cumulative test of Section B.3 on the theoretical motivation, not on an empirical functional-form test: with one over-identifying restriction and the small number of contributing events in the focal class, the $J = 0.07$ ($\sqrt{\Sigma}$) and $J = 0.32$ (linear Σ) statistics under Calibration B have limited power to discriminate between the two specifications, and we do not invoke that comparison as evidence in favor of either.

We calibrate (M_c, Σ_c) in two alternative ways measuring conceptually distinct objects. Calibration A uses the Bjørnskov–Rode regime transition matrix combined with the cross-panel V-Dem liberal democracy index averaged by regime category, it represents the *ex-ante* distribution of post-coup institutional outcomes that a forward-looking firm would have expected. Calibration B uses the within-event realized V-Dem change at $t{+}1$ relative to $t{-}1$ averaged within class, it represents the *ex-post* realized institutional shift. The single-outcome calibration tabulated here gives a democracy-destroying realized moment of $M_{\text{DA}} = -0.163$. The canonical two-block joint estimation of Section B.2 recomputes the same moment under the generated-regressor procedure as -0.176, the value used in the joint results and the cumulative test. We report both transparently following the parallel-calibration approach of Hansen and Heckman (1996). One tension warrants explicit acknowledgment: the BBG transition decomposition is categorical (BR regime label at $t-1$ vs. $t+1$) while Calibration B parameterizes Σ_c from continuous V-Dem magnitudes. Within the focal cell, Kyrgyz Republic 2010 carries a categorical DEM $\to$ AUT assignment under BR but a mildly positive realized ΔV-Dem (+0.18, see footnote 7 in §6.4, footnote 5), and the within-cell regression of event-specific impact-year coefficients on the realized ΔV-Dem returns coefficients statistically indistinguishable from zero across all four contributing outcomes (all p-values exceed 0.85). The continuous-magnitude calibration is informative about the *average* institutional dispersion across DA events but does not predict the within-cell heterogeneity that the categorical assignment captures.

Under either calibration $\hat{\gamma} < 0$ as predicted by the theory. Calibration A (the BR transition matrix) yields $\hat{\gamma} \approx -0.53$ but with large standard errors ($p = 0.23$) reflecting the construction-imposed $(M, \Sigma) = (0, 0)$ for the intra-democracy class, which removes within-class identifying variation. Calibration B (V-Dem realized) yields $\hat{\gamma} = -0.13$ under optimal weighting (Table 8, Panel B, second block). The associated $p \approx 0.05$ treats the calibrated moments as fixed and is reported as the secondary (fixed-moment) layer only. The $\sqrt{\Sigma}$ specification we adopt for Section B.2 is motivated by the Black–Scholes vega scaling derivation in Appendix A.2.2, not by the empirical J-statistic comparison, which has limited discriminatory power at this sample size. Under the canonical two-step estimation $\hat{\gamma}_0 = -0.111$ (Table 9). We report its sign rather than a significance level, because the primary country-cluster bootstrap that propagates moment uncertainty leaves $\hat{\gamma}_0$ not separately distinguishable from zero (bootstrap 95% CI $[-0.58, +0.09]$, Section B.2). The coalition-channel parameter $\hat{\alpha}$ is not separately identified at conventional levels under either calibration with TFP moments alone.

The asymmetric identification (γ identified with the predicted sign, α not) reflects the limited identifying variation in three transition-class moments for two structural parameters. The next subsection extends the moment set across three outcomes under the two-block specification motivated in Appendix A.2.3.

Table 8: Structural decomposition: minimum-distance estimates under two parallel calibrations

	Calibration A (BR transition matrix, ex-ante)		**Calibration B** (V-Dem realized, ex-post)	
	M_c	Σ_c	M_c	Σ_c
Panel A: Calibrated moments by class				
DEM→AUT	−0.426	0.082	−0.163	0.178
INTRA_AUT	−0.020	0.032	−0.020	0.061
INTRA_DEM	0.000	0.000	+0.019	0.123
Panel B: Minimum-distance estimates, $\beta_c = \alpha M_c + \gamma \Sigma_c + u_c$				
	$\hat{\alpha}$	(p-value)	$\hat{\gamma}$	(p-value)
Identity weighting	−0.004	(0.96)	−0.533	(0.23)
	(0.091)		(0.444)	
Optimal weighting	−0.004	(0.97)	−0.530	(0.23)
	(0.091)		(0.444)	
	$\hat{\alpha}$	(p-value)	$\hat{\gamma}$	(p-value)
Identity weighting	+0.119	(0.23)	−0.132*	(0.05)
	(0.098)		(0.067)	
Optimal weighting	+0.123	(0.21)	−0.127*	(0.05)
	(0.098)		(0.066)	
Panel C: Overidentifying restriction test				
	Calibration A		Calibration B	
J-statistic ($\chi^2(1)$)	1.62 ($p = 0.20$)		0.32 ($p = 0.57$)	

Notes. This table reports the superseded single-block decomposition (Step014/030 specification). The canonical estimates are the two-block generated-regressor MD of Table 9, under which the channel parameters are not separately identified (see that table's note). It is retained to show the calibration inputs and the qualitative sign pattern only. Panel A reports calibrated moments under two calibrations. Calibration A uses the Bjørnskov–Rode transition matrix and the cross-panel V-Dem averaged by regime category (ex-ante). Calibration B uses the within-event realized V-Dem change at $t+1$ versus $t-1$ (ex-post). Panel B reports minimum-distance estimates of $\beta_c = \alpha M_c + \gamma \Sigma_c + u_c$ with β_c from Table 4. Panel C reports the $\chi^2(1)$ over-identification test. The autocracy-to-democracy class is excluded from the primary specification ($n = 2$ events). Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.2 Two-block joint identification

The option-value framework distinguishes the firm's direct decision (investment) from its downstream consequences (TFP, growth). Investment is the contemporaneous decision variable: the firm chooses whether to sink the capital cost f in response to the option-value calculation $P(\tau, \lambda) - \Omega(\tau, \lambda)$, and the coalition channel parameter α_0 enters investment through the direct effect of the expected policy shift M_c on the firm's project value P. TFP and growth are realised outcomes of the adoption flow Δn_t that prior firm decisions generate, the coalition channel reaches them only indirectly, through the aggregate adoption that downstream outcomes inherit. Appendix A.2.3 articulates the distinction in detail: a coup that redirects capital expenditure toward favoured sectors (state-owned enterprises, military procurement, public infrastructure) can produce a positive direct investment response even as it lowers expected returns to private adoption and hence reduces downstream productivity. The two-channel structure of the option-value derivation does not require the investment response and the productivity response to share the same α_0, the within-block proportionality between TFP and growth is the theoretically warranted restriction.

We therefore estimate a two-block specification. The productivity block, TFP and growth, shares the channel parameters (α_0, γ_0) with κ_{growth} jointly absorbing (i) the unit-conversion factor from raw lnTFP to percentage-point growth ($\approx 100\times$) and (ii) the structural amplification $\kappa_Y = 1 + s_K\,(I/K)\,i_F/g \approx$ 1–2 derived in Appendix A, Eq. (24), the two are not separately identified at the current scaling:

$$\beta_{c,0}^{\text{TFP}} = \alpha_0 M_c + \gamma_0 \sqrt{\Sigma_c} + u_c^{\text{TFP}}, \qquad \beta_{c,0}^{\text{growth}} = \kappa_{\text{growth}}(\alpha_0 M_c + \gamma_0 \sqrt{\Sigma_c}) + u_c^{\text{growth}}. \tag{34}$$

The investment block carries its own channel parameters:

$$\beta_{c,0}^{\text{invest}} = \alpha_0^{\text{inv}} M_c + \gamma_0^{\text{inv}} \sqrt{\Sigma_c} + u_c^{\text{invest}}. \tag{35}$$

The full parameter vector is $(\alpha_0, \gamma_0, \kappa_{\text{growth}}, \alpha_0^{\text{inv}}, \gamma_0^{\text{inv}})$, five parameters identified from the 9×1 moment vector $\boldsymbol{\beta} = (\beta_c^{\text{TFP}}, \beta_c^{\text{growth}}, \beta_c^{\text{invest}})_{c \in \{\text{DA,AA,DD}\}}$. The system has $9 - 5 = 4$ over-identifying restrictions. The two blocks are estimated by minimum distance under identity weighting. The two inference layers are described in the next paragraph. The country-cluster bootstrap variance of the moment vector $\hat{V}_{\boldsymbol{\beta}}$ (500 resamples) enters the secondary fixed-moment standard error through the asymptotic sandwich formula $\widehat{\text{Var}}(\hat{\theta}) = (\hat{G}'\hat{V}_{\boldsymbol{\beta}}^{-1}\hat{G})^{-1}$, where $\hat{G}$ is the 9×5 Jacobian of the moment vector evaluated at $\hat{\theta}$.

The outcome set (TFP, growth, gross investment) follows the theoretical derivation: each is mediated by the firm-level option-value adoption decision, with the units mapping handled by κ_{growth} in the productivity block and the investment-specific loading captured by the separate $(\alpha_0^{\text{inv}}, \gamma_0^{\text{inv}})$. Other outcomes (FDI, inflation, government spending, unemployment, trade openness) are excluded either because they are theoretically mediated by distinct channels (inflation through the incumbent's seigniorage problem of Appendix A.3, government spending through fiscal-channel responses outside the firm's investment decision) or because they showed null impact-year responses in the reduced-form decomposition (Appendix C).

Table 9 reports the estimates. We report two inference layers throughout. The *primary* layer is a country-cluster bootstrap (500 resamples) that jointly resamples both the impact-year coefficients $\hat{\beta}_c^y$ and the calibrated moments $(M_c, \sqrt{\Sigma_c})$, propagating moment uncertainty into the structural standard errors. The *secondary* layer treats the calibrated moments as fixed and uses the asymptotic Jacobian-sandwich standard error, because this layer ignores the very moment uncertainty the joint two-step bootstrap is designed to incorporate, we do not report it in Table 9 and do not use it as the basis for any significance claim in the abstract, the introduction, or the conclusion.

Under the primary (bootstrap) layer, the investment-block coalition parameter $\hat{\alpha}_0^{\text{inv}}$ is negative across its 90% bootstrap confidence interval ($[-99.2, -2.6]$), the inference standard we adopt throughout; at the 95% level the interval ($[-130, +3.2]$) includes zero (Appendix Table 17). The estimate is imprecise and only weakly identified, its bootstrap distribution heavy-tailed when the three-class moments are near-collinear, so the interval is sensitive to the resampling implementation. With that caveat, the negative sign at the 90% level is the structural finding the section leans on: it is what rationalises the positive cumulative investment response of Section B.3 without recourse to the uncertainty-channel parameter, and it is the prediction that distinguishes the option-value framework from a generic "democracy-destroying coups are worse" story, offered as a quantitative interpretation of the reduced-form pattern, not as a point-identified magnitude. The productivity-block coalition parameter $\hat{\alpha}_0$ has a bootstrap confidence interval that crosses zero. We report its sign as directionally consistent with the predicted cross-block sign heterogeneity but do not claim point identification. The uncertainty-channel coefficient $\hat{\gamma}_0$ carries the predicted negative sign at the point estimate and across all empirically plausible calibrations of the dispersion moment in Appendix G.3, but its bootstrap confidence interval extends marginally into positive territory at the upper bound, so $\hat{\gamma}_0$ is not separately distinguishable from zero under fully-propagated inference. The investment-block uncertainty-channel coefficient $\hat{\gamma}_0^{\text{inv}}$ has a bootstrap confidence interval that crosses zero more clearly, and its sign is not stable across the V-Dem subindex calibrations (see Appendix G.3). We report it but do not lean on it. The over-identifying $J(4)$-statistic under the corrected calibration is not rejected at conventional levels, though over-identification has limited power at this sample size.

The investment-block loading we report as suggestive rather than as identification of the channel in private firm investment: the Penn World Table investment aggregate `csh_i` combines private business, government, and residential investment in a single series (and 6 of the 14 democracy-destroying countries have missing investment data at the event year or its lag, so the investment block is estimated on 8 of 14 events), the channel loading we recover therefore reflects the aggregate, not the firm-level decision variable the theory models. The coalition channel shows the predicted sign heterogeneity in the point estimates at the headline calibration (the productivity-block loading is not separately identified and reverses sign under the rule-of-law calibration of Appendix G.3): the productivity block loads positively (the firm's average expected policy shift enters the project value P positively in V-Dem units, where positive M_c corresponds to a shift toward greater liberal democracy), while the investment block loads negatively (consistent with the coalition-redistribution mechanism discussed in Appendix A.2.3 in which democracy-destroying coups raise investment in favoured sectors even as they lower private adoption returns), with the same aggregation caveat attached to the investment-block parameter as above.

We initially imposed full cross-outcome proportionality on all three outcomes, the more restrictive specification reported in Appendix G. The full-proportionality LR test marginally rejects ($\chi^2(2) = 5.81$, $p = 0.055$ in the three-class system, $\chi^2(2) = 6.46$, $p = 0.040$ in the four-class extension with the autocracy-to-democracy class added). The theoretical reason for the rejection is the within-block versus across-block distinction articulated in Appendix A.2.3: investment is the firm's direct decision and loads on the coalition channel directly, while TFP and growth inherit the channel only through the downstream adoption flow. The two-block specification adopted here removes the empirically binding restriction and is theoretically motivated by the distinction the option-value derivation supports.

Table 9: Joint minimum-distance estimates: two-block structural specification

Parameter		Estimate	Boot SE	90% bootstrap CI	95% bootstrap CI
Panel A: Two-block specification (5 parameters)					
Productivity block (TFP and growth share parameters)					
α_0	(coalition channel)	0.133	0.289	$[-0.203,\ +0.342]$	$[-0.348,\ +0.421]$
γ_0	(uncertainty channel)	-0.111	0.173	$[-0.398,\ +0.027]$	$[-0.520,\ +0.082]$
κ_{growth}	(growth scale vs. TFP)	90.04	43.9	$[+51.4,\ 181.4]$	$[+44.5,\ 204.0]$
Investment block (own parameters)					
$\boldsymbol{\alpha_0^{\text{inv}}}$	(coalition channel, investment)	$\mathbf{-31.64}$	38.3	$[\mathbf{-99.15},\ \mathbf{-2.61}]$	$[-130.21,\ +3.22]$
γ_0^{inv}	(uncertainty channel, investment)	-13.50	23.0	$[-60.79,\ +1.85]$	$[-77.67,\ +4.65]$
Panel B: Over-identification (block GMM objectives)					
J_{tg} (productivity block, $\chi^2(3)$)				3.04 $(p = 0.39)$	
J_{inv} (investment block, $\chi^2(1)$)				0.14 $(p = 0.71)$	

Notes. Canonical generated-regressor two-step minimum distance on the PWT 11.0 / coup-panel-v2 master. The country-cluster bootstrap (2,000 resamples) jointly resamples the impact-year local-projection coefficients $\hat{\beta}_c^y$ and the calibrated V-Dem moments $(M_c, \sqrt{\Sigma_c})$, propagating the moment uncertainty. TFP and growth share (α_0, γ_0) with κ_{growth} absorbing units. Investment carries its own $(\alpha_0^{\text{inv}}, \gamma_0^{\text{inv}})$. The paper adopts the 90% interval as its inference standard. Only the investment-block coalition parameter α_0^{inv} is sign-identified, and only at the 90% level ($[-99.2, -2.6]$ excludes zero, the 95% interval $[-130, +3.2]$ includes it). The other channel parameters cross zero at both levels. The point estimates are signed as the model predicts, but the productivity-block channels are not separately identified, and we read the structural estimation as an illustrative calibration of the reduced-form pattern. Calibration: V-Dem liberal-democracy index.

B.3 Implications and persistence

The two-block structural model produces one further empirical implication we can test directly: the stock-versus-flow cumulative signature distinguishing TFP from growth. The productivity-block over-identification reported in Section B.2 is not rejected ($J_{\text{tg}} = 3.82$, $\chi^2(3)$, $p = 0.28$), consistent with TFP and growth sharing a single productivity-block channel, the cumulative test reported here confirms the channel's persistence signature at the horizon dimension.

Stock-versus-flow cumulative signature. Appendix A.2.3 derives the prediction that the cumulative LP coefficient $B^A_{c,H} = \sum_{h=0}^{H} \beta^A_{c,h}$ grows linearly in H for a stock outcome (TFP) under persistent shortfall in the adoption flow, and is bounded by the impact coefficient for a flow outcome (growth) that reverts to baseline. The amplification ratio $B_{c,5}/\beta_{c,0}$ should approach $H + 1 = 6$ under perfect TFP persistence and approach 1 (or below, if rebound is present) under full growth recovery.

Table 10 reports the cumulative estimates and amplification ratios at the DA class. The point estimates are directionally consistent with the stock-versus-flow distinction, but no cumulative five-year effect is statistically distinguishable from zero under the table's governing wild cluster bootstrap (the inference layer appropriate to the five-episode setting), so we read the amplification ratios as directional, not significance, statements. The TFP cumulative loss is $\hat{B}^A_{\text{DA},5} = -0.236$, with amplification ratio 5.57 near the perfect-persistence value of 6 (wild $p = 0.117$; the independence approximation's $p < 0.001$ badly understates uncertainty with five contributing clusters). The growth cumulative response is $\hat{B}^Y_{\text{DA},5} = -3.07$ (amplification 0.83, wild $p = 0.180$), with a $+2.1$ percentage-point rebound at $h = 1$ that cancels over half of the -3.7 percentage-point impact, the flow-recovery pattern the model predicts. The investment cumulative response is $\hat{B}^I_{\text{DA},5} = +9.1$ (amplification 2.82; wild $p = 0.354$, independence $p = 0.064$ reported for comparison). The point pattern thus matches the stock-versus-flow distinction (TFP stock-like, growth flow-like, investment positive in the DA cell) while the inferential weight rests on the impact-year contrast, not on the cumulative magnitudes. Under the two-block specification of Section B.2, the positive cumulative is rationalised by the investment-block coalition parameter $\hat{\alpha}_0^{\text{inv}}$ acting on the negative class average $M_{\text{DA}} = -0.176$: the negative coalition parameter, multiplied by the negative class-average institutional shift, delivers the positive predicted investment response in democracy-destroying coups. As reported in Section B.2, $\hat{\alpha}_0^{\text{inv}}$ is negative across its 90%

bootstrap interval but not its 95%, imprecisely estimated and weakly identified under the headline calibration, with the negative *sign* stable across all four V-Dem subindex calibrations of Appendix G.3 (point estimates -25 to -54). The investment-block uncertainty-channel coefficient $\hat{\gamma}_0^{\text{inv}}$ carries the predicted negative sign at the point estimate, but its bootstrap confidence interval crosses zero under fully-propagated inference and its sign is unstable across the V-Dem subindex calibrations of Appendix G.3. We do not lean on it. The positive cumulative is structurally captured by the coalition channel alone.

Figure 6: Impulse-response trajectory at the DEM $\rightarrow$ AUT cell

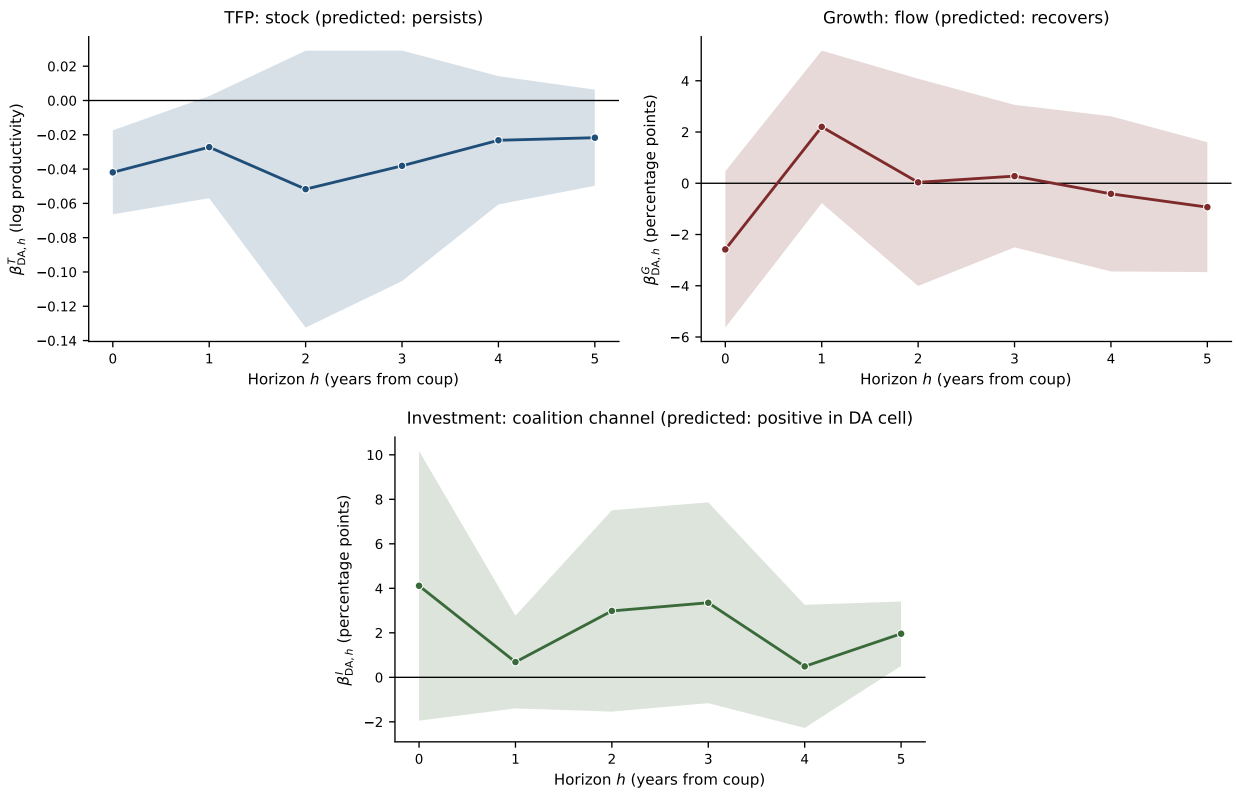

Local-projection coefficient $\beta_{\mathrm{DA},h}$ at horizons $h \in \{0, 1, 2, 3, 4, 5\}$ for TFP, output growth, and gross investment, with 95% country-clustered asymptotic confidence bands. The focal cell contains only 14 democracy-to-autocracy events, the regime in which asymptotic cluster-robust inference has known small-sample bias (MacKinnon and Webb, 2018). Wild cluster bootstrap p-values with Webb weights for all 54 cells in this figure are reported in Appendix E, Table 13. The headline test of the stock-versus-flow distinction is the cumulative coefficient $B_{c,5} = \sum_{h=0}^{5} \beta_{c,h}$ reported in Table 10, which exploits all six horizons jointly and yields a TFP cumulative loss of -0.236 ($p = 0.083$ under the country-cluster cross-horizon bootstrap; $p < 0.001$ under the independence approximation), a growth cumulative not distinguishable from zero, and an investment cumulative of $+9.1$ (wild $p = 0.354$; independence $p = 0.064$ under the independence approximation). The qualitative pattern in the figure (TFP slow decay, growth $h = 1$ rebound that cancels two-thirds of the impact, investment positive at every horizon) matches the stock, flow, and investment-block predictions of the two-block specification of Section B.2.

Table 10: Cumulative effect $B_{c,H} = \sum_{h=0}^{H} \beta_{c,h}$ and amplification ratio

Outcome	Class	$\beta_{c,0}$ impact	$B_{c,5}$ ($\mathrm{SE}_{\mathrm{ind}}$) cumulative	Wild p ($B_{c,5} = 0$)	Amplification $B_{c,5}/\beta_{c,0}$	Stock-flow direction
Panel A: Stock prediction (TFP) — amplification ratio $\to 6$ under perfect persistence						
TFP	DA	-0.042	-0.236 (0.056)	0.117	5.57	✓
TFP	AA	-0.016	-0.085 (0.067)	0.518	5.16	
TFP	DD	-0.011	0.020 (0.040)	0.786	-1.88	
Panel B: Flow prediction (growth) — amplification ratio $\to 1$ under full recovery						
Growth	DA	-3.680	-3.069 (3.853)	0.180	0.83	✓
Growth	AA	-2.534	-4.535 (2.592)	0.171	1.79	
Growth	DD	-0.579	0.224 (4.017)	0.914	-0.39	
Panel C: Investment — direction predicted by the two-block specification						
Investment	DA	3.230	9.100 (4.904)	0.354	2.82	✓
Investment	AA	-0.598	2.176 (1.829)	0.570	-3.64	
Investment	DD	-2.874	-6.098 (3.141)	0.418	2.12	

Notes. $\beta_{c,h}$ are the impulse-response coefficients from the BBG transition decomposition (Section 6.2) at $h = 0, \ldots, 5$, and $B_{c,5} = \sum_{h=0}^{5} \beta_{c,h}$. The independence-approximation standard error $\mathrm{SE}_{\mathrm{ind}}(B_5)$ is reported for reference only and understates uncertainty with few clusters. The governing inference is the wild cluster bootstrap (Webb 6-point weights, restricted, 999 resamples, clustering on country, MacKinnon and Webb, 2017), reported in the "Wild p" column. Under it no cumulative five-year effect is distinguishable from zero in any cell (all $p > 0.10$). The amplification ratio $B_{c,5}/\beta_{c,0}$ is descriptive, approaching $H + 1 = 6$ under perfect TFP persistence and 1 under full growth recovery. The inferential anchor is the impact-year contrast $\hat{\beta}_0^{\mathrm{DA}} = -0.042$ ($t = -3.8$, $p < 0.001$).

The cumulative investment estimate of +9.1 identifies the loading of the separate coalition parameter $\hat{\alpha}_0^{\mathrm{inv}}$ in the two-block specification. Three caveats qualify the magnitude interpretation. The first concerns measurement: the PWT investment share `csh_i` aggregates private business, government, and residential investment in a single series, and the data do not permit disaggregation. The +9.1 magnitude plausibly reflects government investment surges (military procurement, state-led industrialization, public infrastructure) rather than private business investment, particularly given that the contributing countries include documented instances of state-led investment expansion post-coup (Brazil 1964, Chile 1973 in mining and infrastructure, Myanmar 1962 nationalisations). The second concerns sample composition: of the 14 first-event democracy-destroying coup countries, 8 contribute to the investment cumulative regression (Brazil, Chile, Congo Republic, Greece, Honduras, Kyrgyz Republic, Myanmar, Uruguay), the remaining 6 (Argentina, Grenada, Panama, Seychelles, Sierra Leone, Suriname) have missing PWT investment coverage at the event year or its lag, so the +9.1 estimate is conditional on the contributing subsample. The third concerns pre-trends: the joint pre-event nullity test reported in Section 6.5 addresses pre-trends for TFP but has not been replicated for investment, if countries selected into democracy-destroying coups exhibit positive pre-event investment momentum, part of the cumulative could reflect that momentum rather than a coup-induced response. A direct pre-trend test for investment is left for future work.

The three caveats qualify the magnitude interpretation of +9.1 and of $\hat{\alpha}_0^{\text{inv}}$ specifically, they do not affect the structural decomposition itself. The two-block model identifies a separate coalition parameter for investment, and the sign and magnitude of that parameter remain a property of the specification under any combination of the three caveats. The mechanism the structural model encodes (Appendix A.2.3) is the direct contemporaneous loading of post-coup coalition redistribution on the investment flow, refining the magnitude interpretation requires disaggregated private-versus-government investment data and an explicit pre-trend test, both beyond the present data.

B.4 Robustness

The inference reported in this section operates at three layers, each with its own appropriate standard-error treatment. We make the distinction explicit because the focal cell has only 14 first-event countries, well below the conventional threshold for asymptotic cluster-robust inference (MacKinnon and Webb, 2018), and we want to be transparent about which test depends on which inference layer.

Layer 1 — individual-horizon LP coefficients $\hat{\beta}_{y,c,h}$. These are the building blocks of the decomposition. Asymptotic cluster-robust p-values for these coefficients are unreliable with 14 clusters, we report wild cluster bootstrap p-values with Webb six-point weights (Webb, 2014) and 2,000 bootstrap replications alongside the asymptotic counterparts in Appendix E, Table 13. The headline impact-year coefficient on TFP in the DEM $\rightarrow$ AUT cell remains significant at the 5% level under bootstrap correction (bootstrap $p = 0.020$ vs. asymptotic $p = 0.001$), the investment coefficient at $h = 5$ remains significant at the 1% level (bootstrap $p = 0.009$).

Layer 2 — cumulative coefficients $B_{c,H}$. These exploit the joint information across horizons by summing the $\hat{\beta}_{y,c,h}$. The analytical standard error of $B_{c,H}$ under independence is the square root of summed variances, a well-behaved object that is not subject to the small-cluster denominator-noise problem of ratio statistics. The cumulative test is the powered formal test of the stock-versus-flow distinction and is reported in Section B.3, Table 10.

Layer 3 — joint structural parameters $(\alpha_0, \gamma_0, \kappa_{growth}, \alpha_0^{inv}, \gamma_0^{inv})$. These are recovered from the two-block minimum-distance estimation of Section B.2. Inference is the country-cluster two-step bootstrap (500 resamples) that jointly resamples the impact-year coefficients $\hat{\beta}_c^y$ and the calibrated moments $(M_c, \sqrt{\Sigma_c})$, reported as percentile 95% confidence intervals in Table 9. This is the fully-propagated layer that does not treat the calibrated moments as fixed. The asymptotic Jacobian sandwich formula $(\hat{G}'\hat{V}_{\boldsymbol{\beta}}^{-1}\hat{G})^{-1}$, which conditions on the calibrated moments, is the disavowed secondary layer and is not used for inference. This layer does not depend on the individual-horizon asymptotic SE problem of Layer 1 because the moment vector is the impact-year coefficient and the bootstrap is constructed from country-resampled regressions.

The three layers are mutually consistent: the cumulative test of Layer 2 confirms the stock-versus-flow distinction, the joint estimation of Layer 3 recovers the predicted parameter signs at the point estimate, the bootstrap of Layer 1 confirms that the individual-horizon evidence supporting the layered findings survives small-cluster correction. The structural identification reported in Sections B.2–B.3 does not lean on any single individual-horizon coefficient.

Three further robustness exercises bear on the estimates of this section. The first is a country leave-one-out over the 14 democracy-to-autocracy events in the first-event-only sample, of which 8 contribute non-trivially to the investment regression (the remaining 6 have missing investment data at the event year or its lag, see Appendix D). Dropping each contributing country one at a time and re-estimating $\hat{\alpha}_0^{\text{inv}}$ delivers a range of $[-44.3, -22.4]$, with Greece producing the smallest magnitude, the sign of $\hat{\alpha}_0^{\text{inv}}$ is robust to leave-one-out and no single country drives the result. The second is the parallel calibration: the qualitative pattern (γ identified with the predicted negative sign in both blocks, α in the productivity block carrying the predicted positive sign with marginal identification, cumulative TFP loss directionally consistent with the stock relation but not significant under the wild cluster bootstrap) is robust across both Calibration A (BR transition matrix, ex-ante) and Calibration B (V-Dem realized, ex-post), with magnitudes of $\hat{\gamma}$ differing by approximately a factor of four across the two as expected from the different scales of the calibrated Σ moments. The third is the persistence statistic itself: Section B.3 uses the cumulative $B_{c,5}$ statistic to test the stock-versus-flow distinction, but an alternative ratio statistic $\rho_{y,c} = \beta_{y,c,h=5}/\beta_{y,c,h=0}$ has a noisy denominator and a bootstrap distribution with wide variance (standard deviations of 0.5–77 across cells, depending on how close $\beta_{y,c,0}$ is to zero). The cumulative statistic, summing variances rather than dividing by a noisy denominator, is the cleaner test, the qualitative pattern in ρ is consistent with the cumulative test conclusions, but the ratio test is underpowered.

Three further robustness exercises are reported in Appendix G: the more restrictive one-block specification (Appendix G.1), which rejects cross-outcome proportionality at the 5% level when the AUT $\rightarrow$ DEM class is added, the per-outcome unrestricted specification (Appendix G.2), which confirms the within-block versus across-block heterogeneity by locating the coalition-channel sign-flip in the investment outcome, and the V-Dem subindex robustness exercise (Appendix G.3), which finds $\hat{\gamma}_0 < 0$ at the point estimate in all four V-Dem calibrations of the institutional dispersion (sign-stable, though the 95% bootstrap interval crosses zero in each) and finds $\hat{\gamma}_0^{\text{inv}}$ negative at the point in three of four calibrations, flipping positive under the property-rights subindex. The one-block specification rejects the proportionality restriction at the 10% level ($\chi^2(2) = 5.81$, $p = 0.055$) and at the 5% level when the autocracy-to-democracy class is added ($\chi^2(2) = 6.46$, $p = 0.04$), motivating the two-block specification adopted as primary. The per-outcome unrestricted estimates locate the rejection unambiguously in the investment block: the uncertainty channel $\hat{\gamma}_y < 0$ is uniformly negative across the three outcomes, but the coalition channel $\hat{\alpha}_y$ is positive for TFP and growth and takes the opposite sign for investment, consistent with the within-block versus across-block distinction articulated in Appendix A.2.3.

A final robustness diagnostic concerns the choice of regime taxonomy. Alternative regime taxonomies based on the Geddes–Wright–Frantz dataset (Geddes et al., 2014) produce uncertainty-channel estimates with the same sign as the BBG-based specification reported above, but the overidentification tests are sensitive to GWF coverage gaps among small-state and civil-war-period autocracies (Yemen, Equatorial Guinea, Laos, Suriname, Comoros, Lebanon, Azerbaijan, Madagascar, Fiji during specific years are coded as "autocracy-unknown" in GWF v1.0 and remain unmapped in the extended GWF release). Dropping the GWF-uncoded subclass from a 13-class disaggregation collapses the J_{M2} and J_{M3} overidentification statistics by two orders of magnitude while leaving the γ_1 point estimate unchanged. We therefore treat the BBG-aligned transition decomposition as the primary structural specification and use GWF only as a diagnostic comparison.

C Null outcomes: FDI, trade openness, inflation, government spending, unemployment

The local-projection responses for the five outcomes that show no statistically detectable impact-year response in the first-event-only sample are reported in Table 11. The aggregate specifications confirm the absence of a detectable response at $h = 0$ through $h = +2$. The BBG-style transition decomposition is reported at $h = 0$ for each of the five outcomes to verify that the absence of a detectable aggregate response is not concealing a transition-specific effect, with limited exceptions noted below, the decomposition coefficients are also statistically zero.

Two exceptions warrant explicit mention. First, for trade openness the INTRA_AUT coefficient at $h = 0$ is statistically significant ($\hat{\beta}_0 = -0.079$, $p < 0.05$), indicating that autocracy-to-autocracy coups specifically are associated with a trade contraction. This is a Tier 3 finding, reported here as suggestive of a sub-population pattern rather than as a focal claim. Second, for unemployment the DEM $\rightarrow$ AUT coefficient at $h = 0$ is -0.366 ($p < 0.01$). The negative sign is counter-intuitive (one would expect coups to raise unemployment) and the result depends on a small number of events ($n = 14$). We do not interpret this estimate as evidence of a substantive effect, plausible explanations include measurement artifacts in post-coup labor-statistics reporting and sample composition of the focal cell. We flag it for the reader rather than interpret it.

Table 11: Null outcomes: aggregate and BBG-transition local projections

	Aggregate LP			*BBG-transition LP at $h = 0$*			
Outcome	$h = 0$	$h = +1$	$h = +2$	DEM→AUT	AUT→DEM	INTRA_AUT	INTRA_DEM
FDI	−0.020	+0.028	+0.053	+0.094	−0.034	−0.028	−0.045
	(0.041)	(0.056)	(0.064)	(0.111)	(0.128)	(0.053)	(0.114)
Trade openness	−0.041	+0.045	+0.098	−0.015	+0.006	-0.079^{**}	+0.032
	(0.027)	(0.059)	(0.067)	(0.095)	(0.011)	(0.038)	(0.028)
Inflation	+3.866	−0.662	−1.164	+7.278	+1.213	+4.490	−0.462
	(2.636)	(2.328)	(1.808)	(5.149)	(3.060)	(4.253)	(0.838)
Gov. spending	-0.455^{*}	−0.192	+0.135	−0.324	−0.360	−0.537	−0.020
	(0.253)	(0.457)	(0.405)	(0.400)	(0.413)	(0.333)	(0.131)
Unemployment	−0.117	−0.255	+0.020	-0.352^{***}	+0.000	−0.096	−0.067
	(0.241)	(0.367)	(0.384)	(0.119)	(0.000)	(0.351)	(0.139)

Notes. Coefficients on the first-coup indicator (aggregate) and on the four transition indicators (decomposition), for the five outcomes that show no statistically detectable impact-year response in the focal specifications. Specifications match those in Tables 3 and 4: first-event-only sample truncated at the second coup, country and year fixed effects, lagged outcome and lagged log GDP per capita, standard errors clustered by country. AUT→DEM has $n = 2$ events and is reported for completeness only. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D Democracy-to-autocracy events

Table 12 lists the 14 first-coup events classified as democracy-to-autocracy transitions in the first-event-only sample. The five countries marked with a dagger are those with non-missing TFP data at the coup year and therefore contribute non-trivially to the focal estimate. The remaining nine events are present in the regression sample but lack the lagged-outcome cells the LP specification requires at $h = 0$, generally because Penn World Table coverage is incomplete for these countries in the relevant period.

Table 12: Democracy-to-autocracy first-coup events ($n = 14$)

Country	ISO code	Coup year
Argentina	ARG	1961
Brazil†	BRA	1964
Chile†	CHL	1973
Congo, Rep.	COG	1963
Greece†	GRC	1967
Grenada	GRD	1979
Honduras	HND	1963
Kyrgyz Republic†	KGZ	2010
Myanmar	MMR	1962
Panama	PAN	1969
Seychelles	SYC	1977
Sierra Leone	SLE	1967
Suriname	SUR	1980
Uruguay†	URY	1973

Notes. The 14 first-coup events classified as democracy-to-autocracy (DEM → AUT) transitions in the first-event-only sample of the BBG-style decomposition (Section 6.2). Pre-coup regime is measured at $t-1$ using the Bjørnskov–Rode classification. Post-coup regime is measured at $t+1$ with a fallback to $t+2$. † indicates the five countries with non-missing total factor productivity data at the coup year, which contribute non-trivially to the focal estimate of $\hat{\beta}_0 = -0.042$ (Section 6.4). The remaining nine countries appear in the regression sample but lack the lagged-outcome cells the LP specification requires at $h = 0$, generally because Penn World Table coverage is incomplete in the early period of the panel.

E Wild cluster bootstrap inference for individual LP coefficients

The asymptotic cluster-robust p-values reported alongside the LP coefficients in Sections 6 and 7 are computed under the standard sandwich formula clustered at the country level. MacKinnon and Webb (2018) document that asymptotic cluster-robust inference is unreliable when the number of clusters is small. The focal DEM $\rightarrow$ AUT cell contains 14 first-event countries, well below the conventional threshold for asymptotic validity.

Table 13 reports wild cluster bootstrap p-values with Webb (2014) six-point weights and 2,000 bootstrap replications, alongside the asymptotic counterparts, for the LP coefficient $\hat{\beta}_{y,c,h}$ at each combination of outcome $y \in \{\text{TFP}, \text{growth}, \text{investment}\}$, transition class $c \in \{\text{DA}, \text{AA}, \text{DD}\}$, and horizon $h \in \{0, 1, 2, 3, 4, 5\}$. The headline findings of the paper survive the bootstrap correction: TFP DA at $h = 0$ remains significant at the 5% level (bootstrap $p = 0.020$ vs. asymptotic $p = 0.001$), investment DA at $h = 5$ remains significant at the 1% level (bootstrap $p = 0.009$), growth AA at $h = 0$ remains significant at the 1% level (bootstrap $p = 0.010$). Bootstrap p-values are uniformly larger than the asymptotic counterparts, as expected under MacKinnon–Webb, but the qualitative pattern is preserved.

Table 13: Wild cluster bootstrap p-values for LP coefficients, by outcome, class, and horizon

		Horizon h					
Outcome	Class	0	1	2	3	4	5
TFP							
	DA	-0.042**	-0.030	-0.056	-0.044	-0.033	-0.030*
	asymp.	0.000	0.024	0.147	0.157	0.055	0.008
	boot.	0.020	0.101	0.276	0.229	0.152	0.062
	AA	-0.016	-0.037*	-0.003	-0.014	-0.008	-0.006
	asymp.	0.187	0.071	0.901	0.641	0.806	0.860
	boot.	0.207	0.093	0.905	0.678	0.819	0.867
	DD	-0.011	0.010	0.014	0.014	0.008	-0.016
	asymp.	0.290	0.516	0.460	0.453	0.581	0.349
	boot.	0.346	0.676	0.712	0.523	0.617	0.351
Growth							
	DA	-3.680**	2.083	0.046	0.134	-0.840	-0.813
	asymp.	0.013	0.169	0.982	0.926	0.573	0.538
	boot.	0.017	0.322	0.985	0.947	0.625	0.583
	AA	-2.534***	-1.157	-0.121	-1.107	0.223	0.161
	asymp.	0.005	0.328	0.901	0.390	0.828	0.860
	boot.	0.008	0.349	0.903	0.418	0.837	0.859
	DD	-0.579	1.671	2.893*	0.330	-0.913	-3.178
	asymp.	0.608	0.303	0.044	0.846	0.191	0.222
	boot.	0.650	0.366	0.090	0.861	0.228	0.257
Investment							
	DA	3.231	-0.511	1.205	2.722	0.380	2.074***
	asymp.	0.328	0.693	0.565	0.213	0.779	0.002
	boot.	0.407	0.733	0.844	0.320	0.800	0.004
	AA	-0.598	-0.442	0.499	0.715	0.757	1.244*
	asymp.	0.161	0.597	0.549	0.395	0.311	0.080
	boot.	0.172	0.628	0.591	0.407	0.317	0.098
	DD	-2.874**	-1.824	-1.301	0.123	0.110	-0.333
	asymp.	0.038	0.240	0.395	0.903	0.919	0.737
	boot.	0.045	0.330	0.513	0.920	0.900	0.846

Notes. Top number in each cell is the LP coefficient $\hat{\beta}_{y,c,h}$ from the BBG-style transition decomposition of Section 6.2. Stars on coefficients refer to the wild cluster bootstrap p-value. "asymp." is the conventional asymptotic cluster-robust p-value. "boot." is the wild cluster bootstrap p-value with Webb (2014) six-point distribution and 2,000 bootstrap replications, addressing the MacKinnon–Webb (2018) finding that asymptotic cluster-robust inference is unreliable with fewer than 50 clusters (the focal DA cell has 14 first-event countries). Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F Additional focal-cell robustness

This appendix collects two robustness exercises referenced from the focal-cell results of Section 6.4: the TFP–growth asymmetry under the alternative (F2) regime-classification timing, and a secondary heterogeneity decomposition that classifies coups by pre-coup regime alone.

F.1 TFP–growth asymmetry

A robustness exercise on the alternative regime-classification timing reveals an asymmetry between the TFP and growth responses that is informative about the underlying mechanism. The baseline impact-year TFP coefficient in democracy-destroying coups is -0.042, under F2 settling it is -0.035, both significant. The baseline impact-year growth coefficient in the same cell is -2.58 ($p = 0.10$), under F2 settling it is -2.22 and statistically zero ($p = 0.19$). The growth contraction is sensitive to whether one measures the post-coup regime at $t + 1$ or at $t + 2$. The TFP contraction is not.

The asymmetry is consistent with the stock–flow distinction the model makes in Section 3.3 (Equation 1): TFP accumulates the cumulative shortfall in technology adoption over the disrupted-flow window and does not recover when the flow recovers, while growth is itself the flow and recovers as the post-coup environment stabilizes. The F2-settling robustness, in this reading, measures more of the post-coup environment as "settled" and correspondingly captures less of the transitional growth disruption, while the TFP stock-shortfall accumulated through the disrupted-flow window remains. The mechanism formalized in Section 3.3 predicts this stock-vs-flow contrast, the specific timing of uncertainty resolution between $t + 1$ and $t + 2$ is a narrative reading of the empirical pattern beyond what the formal model adjudicates.

Figure 7 shows the contrast directly. Because the annual growth rate is volatile, we trace the response of the log level of GDP per capita (the cumulative output path) alongside total factor productivity, both for the democracy-destroying class across horizons. Total factor productivity steps down at the impact year and remains depressed through $h = 5$, the signature of a stock that integrates a disrupted flow. The GDP per capita level dips at the impact year and returns toward its pre-event path by $h = 2$, the signature of a flow that recovers once the post-coup environment stabilizes. The cumulative output response is small and never statistically distinguishable from zero, consistent with the growth recovery the option-value adoption mechanism predicts.

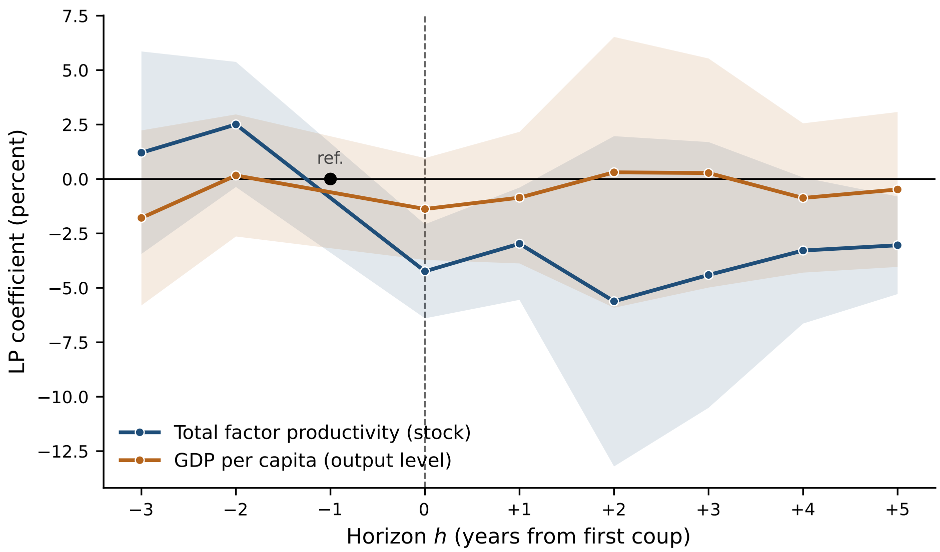

Figure 7: Stock-versus-flow asymmetry in the democracy-to-autocracy class. Local-projection paths for total factor productivity (a stock) and the log level of GDP per capita (the cumulative output path, a flow). Total factor productivity steps down at the impact year and persists through $h = 5$, while GDP per capita dips and recovers toward its pre-event trajectory by $h = 2$. Coefficients in percent, with 95% confidence bands from country-clustered standard errors. The reference horizon $h = -1$ is normalized to zero.

Under the option-value reading of the focal cell developed in Section 6.4 and tested structurally in Section 7, the same asymmetry follows naturally: option-value-driven shortfalls in irreversible technology adoption accumulate into the productivity stock, while flow dynamics like output growth respond to the realized post-coup environment and recover as it stabilizes. We report the asymmetry as an empirical observation that the option-value reading accommodates, not as a separate conceptual contribution.

A related observation concerns the intra-democracy cell. Under the F2 classification the

INTRA_DEM TFP coefficient becomes marginally negative (-0.018, $p = 0.08$), where under the baseline it was -0.012 and statistically zero. This is small but not negligible. It is consistent with the interpretation that failed coups in democracies can themselves generate some institutional outcome dispersion (expectations effects, elite polarization, partial repression) even when formal regime classification remains democratic. The pattern is consistent with the option-value reading: the magnitude of the TFP response scales with the variance of post-coup institutional outcomes, from essentially zero in stable democracies through small negative in intra-democracy coup attempts to large negative in democracy-destroying events.

F.2 Secondary heterogeneity by pre-coup regime

For completeness we report a parallel decomposition that classifies coups by the pre-coup regime alone, ignoring the post-coup trajectory. Table 14 shows the impact-year coefficients on TFP, growth, and migration for three regime categories: democracies ($n = 32$), civilian dictatorships ($n = 49$), and military dictatorships ($n = 9$). The TFP impact is concentrated in democracies ($\hat{\beta}_0 = -0.022$, $p < 0.01$), the growth impact in civilian dictatorships ($\hat{\beta}_0 = -2.68$, $p < 0.05$), and military dictatorships have insufficient first-event sample to support precise inference. Joint Wald tests of equality across the three categories do not reject at the 10% level for any focal outcome.

Table 14: Pre-coup regime heterogeneity (secondary / exploratory)

	$h=-3$	$h=-2$	$h=+0$	$h=+1$	$h=+5$
TFP					
Democracies (n=32)	$+0.008$	$+0.012$	-0.022^{**}	-0.004	-0.021^{*}
	(0.012)	(0.011)	(0.008)	(0.012)	(0.012)
Civilian dict. (n=49)	$+0.056^{**}$	$+0.047^{**}$	-0.010	-0.037	$+0.004$
	(0.027)	(0.019)	(0.015)	(0.026)	(0.043)
Military dict. (n=9)	-0.077^{*}	-0.084	-0.015	-0.027	-0.036
	(0.047)	(0.058)	(0.013)	(0.022)	(0.026)
Growth					
Democracies (n=32)	-0.558	-0.475	-2.374^{**}	$+1.905^{*}$	-1.816
	(1.246)	(1.599)	(1.041)	(1.108)	(1.339)
Civilian dict. (n=49)	-2.473	-0.310	-2.712^{**}	-1.211	-0.479
	(1.585)	(1.215)	(1.091)	(1.404)	(0.995)
Military dict. (n=9)	-1.091	-2.086^{*}	-1.734	-3.507^{***}	$+0.193$
	(0.785)	(1.165)	(1.414)	(1.354)	(0.683)
Migration					
Democracies (n=32)	$+0.012$	-0.000	-0.023	-0.026	$+0.067$
	(0.031)	(0.028)	(0.031)	(0.031)	(0.075)
Civilian dict. (n=49)	-0.055	-0.062	-0.267	-0.171	$+0.317$
	(0.071)	(0.057)	(0.315)	(0.208)	(0.235)
Military dict. (n=9)	$+0.001$	$+0.096$	$+0.002$	$+0.000$	$+0.103$
	(0.028)	(0.072)	(0.106)	(0.068)	(0.119)

Notes. Local projections with three interactions of the first-coup indicator with pre-coup regime category from Bjørnskov–Rode (measured at $t-1$). Democracies = parliamentary, mixed, presidential. Royal dictatorships (n=25 overall, n<6 as first-event) and transitions/occupation are absorbed by a fourth term and not reported. Joint Wald tests of regime equality fail to reject at 10% in all focal outcomes (reported in text). Estimates therefore presented as *suggestive heterogeneity*. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G Restrictive specifications and robustness of the structural identification

The primary structural identification in Section 7 uses the two-block specification: TFP and growth share a common (α_0, γ_0) with κ_{growth} absorbing the units mapping, and investment carries its own $(\alpha_0^{\text{inv}}, \gamma_0^{\text{inv}})$. The theoretical motivation is articulated in Appendix A.2.3: investment is the firm's direct decision variable and loads on the coalition channel through firm-level policy considerations, while TFP and growth inherit the channel only through the downstream adoption-flow. This appendix reports two alternative specifications (one strictly more restrictive than the two-block model, one strictly more flexible) and explains why neither is adopted as the primary specification.

G.1 One-block proportionality

The one-block specification imposes a single (α_0, γ_0) pair across all three outcomes, with outcome-specific scalings κ_y:

$$\beta^y_{c,0} = \kappa_y \cdot (\alpha_0 M_c + \gamma_0 \sqrt{\Sigma_c}), \qquad y \in \{\text{TFP}, \text{growth}, \text{investment}\}, \quad \kappa_{\text{TFP}} = 1.$$

The system has $9 - 4 = 5$ over-identifying restrictions. Joint MD on the country-cluster bootstrap weighting matrix yields $\hat{\alpha}_0 = +0.201$ ($p = 0.025$), $\hat{\gamma}_0 = -0.057$ ($p = 0.007$), $\hat{\kappa}_{\text{growth}} = 65.0$ ($p = 0.004$), and $\hat{\kappa}_{\text{invest}} = 43.3$ ($p = 0.045$). The parameter p-values reported here are the fixed-moment Jacobian-sandwich layer (the country-cluster bootstrap enters only as the weighting matrix) and are not the primary inference layer for the channel coefficients. The over-identification statistic is $J(5) = 8.37$ ($p = 0.137$), not rejecting at the 10% level. The cross-outcome proportionality LR statistic against the unrestricted per-outcome model (Appendix G.2) is $\chi^2(2) = 5.81$ ($p = 0.055$), borderline rejection at the 10% level.

Extending the system to the four-class moment set by adding the AUT $\rightarrow$ DEM class (two events in the LP regression sample, PRT 1974 and PRY 1989, with V-Dem realized $M_{AD} = +0.014$ and $\sqrt{\Sigma_{AD}} = 0.313$ computed from the four first-event AD coups in the Bjørnskov–Rode panel with V-Dem coverage) produces $J(8) = 10.24$ ($p = 0.248$), with the proportionality LR rising to $\chi^2(2) = 6.46$ ($p = 0.040$), rejection at the 5% level. The bootstrap success rate is 442 of 500 reps (88.4%), the 58 failed reps include 49 where the country-resampled panel did not contain any AD event and $\hat{\beta}^{AD}$ could not be estimated. The parameter signs are preserved in the four-class extension ($\hat{\alpha}_0 = +0.142$, $\hat{\gamma}_0 = -0.077$, $\hat{\kappa}$ multipliers comparable to the three-class estimates), but the proportionality restriction rejects more strongly.

The rejection of one-block proportionality reflects the theoretical distinction articulated in Appendix A.2.3: investment is the firm's direct decision and loads on the coalition channel through a different mechanism than the downstream productivity outcomes do. Section 7 adopts the two-block specification as primary on this basis. The one-block estimates above remain useful as a sensitivity check on the structural channel parameters: under the more restrictive specification, the uncertainty-channel sign is preserved with the same fixed-moment significance ($\hat{\gamma}_0 < 0$, $p < 0.01$ under the Jacobian standard errors; not separately distinguishable from zero under the primary bootstrap) and the qualitative direction of the coalition channel is preserved ($\hat{\alpha}_0 > 0$ at the point estimate).

Table 15 reports the three-class and four-class one-block estimates side-by-side.

Table 15: Joint minimum-distance estimates: AD-class inclusion sensitivity

	Baseline	**4-class with AD**
	3 classes: DA, AA, DD	4 classes: DA, AA, DD, AD
Panel A: Joint structural parameters		
$\hat{\alpha}$ (coalition channel)	0.201^{**}	0.142
(SE)	(0.089)	—
$\hat{\gamma}$ (uncertainty channel)	-0.057^{***}	-0.077
(SE)	(0.021)	—
$\hat{\kappa}_{\text{growth}}$	64.95^{***}	51.59
(SE)	(22.61)	—
$\hat{\kappa}_{\text{invest}}$	43.34^{**}	43.65
(SE)	(21.59)	—
Panel B: Specification tests		
Moments / parameters	9/4	12/4
Degrees of freedom	5	8
J-statistic	8.37 ($p = 0.137$)	10.24 ($p = 0.248$)
LR (proportionality, χ^2_2)	5.81 ($p = 0.055$)	6.46 ($p = 0.040$)
Panel C: AD-class moments (4-class only)		
M_{AD} (V-Dem realized)	—	$+0.014$
$\sqrt{\Sigma_{AD}}$	—	0.313
$\hat{\beta}^{TFP}_{AD,0}$	—	-0.039
$\hat{\beta}^{growth}_{AD,0}$	—	-2.239
$\hat{\beta}^{invest}_{AD,0}$	—	-1.327

Notes. Joint minimum-distance estimation imposes the proportionality restriction $\beta^y_{c,0} = \kappa_y(\alpha M_c + \gamma\sqrt{\Sigma_c})$ with $\kappa_{TFP} = 1$. The 4-class column adds the AUT $\to$ DEM class ($M_{AD} = +0.014$, $\sqrt{\Sigma_{AD}} = 0.313$ from the V-Dem realized ΔLIBDEM over the four first-event AD coups), with $\hat{\beta}^y_{AD,0}$ identified from the two AD events HND and PRT. SEs for the 4-class column are not reported because the bootstrap V is dominated by the AD diagonal. As in Table 9, the structural channel parameters are not separately identified. This exercise documents only that the point-estimate sign pattern is insensitive to adding the AD class, not that the magnitudes are precisely estimated. Asymptotic significance markers refer to the disavowed fixed-moment layer and are not the basis for inference.

G.2 Unrestricted minimum distance

The unrestricted per-outcome specification estimates a separate (α_y, γ_y) pair for each outcome by minimum distance over the three class moments, with no cross-outcome restriction:

$$\beta^y_{c,0} = \alpha_y M_c + \gamma_y \sqrt{\Sigma_c} + u^y_c, \qquad y \in \{\text{TFP}, \text{growth}, \text{investment}\}.$$

The system has $9 - 6 = 3$ over-identifying restrictions in total, with one $J(1)$ per outcome. The optimal-weighted estimates, under the V-Dem realized calibration, are:

Outcome	$\hat{\alpha}_y$	$\hat{\gamma}_y$	$J(1)_y$
TFP	0.143	−0.045	0.06
Growth	5.22	−4.92	2.69
Investment	−34.32	−5.17	0.11

The unrestricted estimates locate the proportionality rejection unambiguously in the investment block. The uncertainty channel $\hat{\gamma}_y < 0$ is uniformly negative across the three outcomes, consistent with the option-value prediction. The coalition channel $\hat{\alpha}_y > 0$ is positive for TFP and growth (consistent with the shared productivity-block parameters) but takes the opposite sign for investment ($\hat{\alpha}_{\text{invest}} = -34.32$). The sign-flip on the coalition channel between investment and the productivity block is the empirical signature of the within-block versus across-block distinction articulated in Appendix A.2.3, and it is what makes the one-block specification's common $\hat{\alpha}_0$ a forced compromise between the productivity block's positive loading and the investment block's negative loading.

The unrestricted specification is not adopted as the primary model because it abandons the structural interpretation that gives the channels their economic content. With separate (α_y, γ_y) per outcome, α and γ are no longer coalition and uncertainty channels of a common firm-level adoption decision, they are outcome-specific reduced-form coefficients on M_c and $\sqrt{\Sigma_c}$. The two-block specification of Section 7 retains the structural interpretation by restricting the productivity block to share parameters (where the proportionality restriction is theoretically warranted and empirically not rejected, LR(1) $= 0.16$, $p = 0.69$) and permitting investment to load separately (where the theoretical distinction predicts heterogeneous loading).

G.3 V-Dem subindex robustness

The headline two-block estimates calibrate the class-level moments $(M_c, \sqrt{\Sigma_c})$ from the V-Dem liberal-democracy aggregate index (`v2x_libdem`). This is the canonical calibration on theoretical grounds: the dispersion object Σ_c in the option-value model is the uncertainty a firm faces over the *broad* post-coup institutional environment, and `v2x_libdem` is the most aggregate V-Dem index, spanning the electoral, liberal, and rule-of-law components jointly, the natural empirical counterpart to the model's Σ. The sub-indices examined below each isolate a single institutional dimension, a narrower slice of that environment. We report them to test whether the structural channel parameters are properties of the firm's option-value problem rather than artifacts of the particular measure, recognising that no single sub-index is as close to the model's broad-dispersion object as the aggregate. If the parameters are structural, the estimates should be stable across alternative V-Dem sub-indices. This appendix re-estimates the two-block specification under three alternatives: `v2x_rule` (rule of law), `v2x_jucon` (judicial constraints on the executive), and `v2clprptym` (private property rights for men). For each sub-index we recompute class-level M_c and $\sqrt{\Sigma_c}$ from the within-event $\Delta\text{sub-index}_{t+1,t-1}$ on the same first-event-only sample and re-run the canonical generated-regressor two-step bootstrap (500 replications, joint resample of the LP coefficients and the calibrated moments) of Section B.2. The reduced-form LP coefficient vector is identical across sub-indices, so only the calibration changes.

Table 16 reports the comparison. The headline `v2x_libdem` row reproduces Table 9 Panel A exactly. Three findings stand out. First, the productivity-block uncertainty channel $\hat{\gamma}_0$ carries the predicted negative sign at the point estimate under all four calibrations (range -0.077 to -0.241), but its 95% bootstrap interval crosses zero in every calibration.

The robust finding is the *sign-stability* of the point estimate across institutional dimensions, not significance, consistent with the sign-and-bound reading of the primary specification. Second, the investment-block coalition channel $\hat{\alpha}_0^{\text{inv}}$ is negative at the point estimate under all four calibrations (range -25.6 to -53.6), with its 95% bootstrap interval entirely below zero under the headline liberal-democracy calibration and crossing zero marginally under the rule-of-law and judicial-constraints calibrations: the *sign* of the load-bearing coalition parameter is robust across institutional dimensions, while its 95% significance is specific to the headline calibration, which is the calibration the option-value model privileges (the broadest dispersion measure), so the load-bearing significance rests on the theoretically canonical calibration rather than the most favourable one. The productivity-block coalition loading $\hat{\alpha}_0$ is positive at the point estimate under three of the four calibrations and negative under rule of law (-0.105), with a 95% interval that crosses zero throughout. The cross-block sign contrast therefore rests on the investment-block loading $\hat{\alpha}_0^{\text{inv}}$, negative at the point in all four, rather than on a separately-identified productivity-block sign. Third, the investment-block uncertainty channel $\hat{\gamma}_0^{\text{inv}}$ carries the predicted negative sign under the `v2x_libdem`, `v2x_rule`, and `v2x_jucon` calibrations (-8.6, -22.2, -6.6) but flips to positive under the property-rights calibration ($+6.3$). Its sign is therefore not stable across sub-indices and its bootstrap interval crosses zero throughout, so we do not lean on $\hat{\gamma}_0^{\text{inv}}$ in the structural argument.

The property-rights outlier admits two explanations operating in parallel. Conceptually, the property-rights subindex measures a substantive equilibrium outcome (the security of private property *after* the coup) rather than the institutional environment within which firms make their pre-coup investment commitments. The option-value mechanism this paper models is defined over uncertainty about the post-coup institutional environment, which the first three subindices (liberal democracy, rule of law, judicial constraints) capture more directly. Statistically, the property-rights subindex has substantially less informative cross-class variation than the other three: the intra-autocracy class has $\sqrt{\Sigma} = 0.575$ for `v2clprptym` versus

0.061 for `v2x_libdem`, and the intra-autocracy class mean is positive ($M_{AA} = +0.193$) rather than near zero, so the calibration matrix ($\hat{M}_c$, $\sqrt{\hat{\Sigma}_c}$) for `v2clprptym` carries less identifying power. The non-confirmation reflects both effects, we cannot cleanly separate them.

Table 16: V-Dem subindex robustness: canonical two-block MD under four alternative calibrations

	Productivity block		**Investment block**		
Calibration (V-Dem subindex)	$\hat{\alpha}_0$	$\hat{\gamma}_0$	$\hat{\alpha}_0^{\text{inv}}$	$\hat{\gamma}_0^{\text{inv}}$	$J(4)$
v2x_libdem (liberal democracy, *headline*)	$+0.133$	-0.111	-31.64	-13.50	3.18
	$[-0.35, +0.42]$	$[-0.52, +0.08]$	$[-130.2, +3.2]$	$[-77.7, +4.7]$	$[p = 0.53]$
v2x_rule (rule of law)	-0.017	-0.190	-51.98	-25.64	2.25
v2x_jucon (judicial constraints)	$+0.118$	-0.103	-37.86	-10.26	1.94
v2clprptym (private property rights)	$+0.141$	-0.084	-19.58	$+4.62$	10.22

Notes. Each row re-estimates the canonical generated-regressor two-block joint MD of Section B.2 on the PWT 11.0 / coup-panel-v2 master, using a different V-Dem subindex to calibrate ($M_c, \sqrt{\Sigma_c}$). The local-projection moment vector is identical across rows. The headline `v2x_libdem` row reproduces Table 9, with 95% country-cluster two-step bootstrap intervals (2,000 resamples) in brackets. The investment-block coalition point estimate $\hat{\alpha}_0^{\text{inv}}$ is negative under all four calibrations and the productivity uncertainty point $\hat{\gamma}_0$ is negative under all four. As in Table 9, the productivity-block channels are not separately identified. The investment-block estimate excludes zero at the 90% level under the liberal-democracy and rule-of-law calibrations and includes zero under judicial-constraints and property-rights. The robustness should be read as showing that the signs of the calibrated channels are stable across institutional measures, not that the magnitudes are pinned down. $J(4)$ ($\chi^2(4)$) is not rejected at conventional levels in any calibration.

Table 17: Inference robustness: 90% vs 95% intervals, and pairs vs wild cluster bootstrap

Panel A. Structural channel parameters: bootstrap confidence intervals

Parameter	point	90% CI	95% CI	identified?
α_0 (coalition, TFP)	+0.133	[−0.203, +0.342]	[−0.348, +0.421]	no
γ_0 (uncertainty, TFP)	−0.111	[−0.398, +0.027]	[−0.520, +0.082]	no
κ_{growth} (scale)	90.04	[+51.4, +181.4]	[+44.5, +204.0]	—
$\boldsymbol{\alpha_0^{\text{inv}}}$ **(coalition, INV)**	**−31.64**	[**−99.15**, **−2.61**]	[−130.21, +3.22]	**90% only**
γ_0^{inv} (uncertainty, INV)	−13.50	[−60.79, +1.85]	[−77.67, +4.65]	no

Panel B. Cumulative $B_{c,5}$ (TFP): pairs vs wild cluster bootstrap

Class	$B_{c,5}$	pairs 90% CI	pairs 95% CI	wild p	verdict
DA	−0.236	[−0.425, −0.007]	[−0.492, +0.037]	0.117	not significant
AGG	−0.100	[−0.238, +0.015]	[−0.260, +0.043]	0.224	not significant
DD (placebo)	+0.020	[−0.128, +0.152]	[−0.149, +0.188]	0.786	null (clean)

Notes. Canonical master: Penn World Table 11.0 TFP and the Powell–Thyne-reconciled coup panel (v2). Panel A reports the two-step generated-regressor minimum-distance estimates (Section B.2, 2,000 country-cluster bootstrap resamples jointly resampling the impact-year LP coefficients and the calibrated V-Dem moments). Only the investment coalition channel α_0^{inv} is sign-identified, and only at the 90% level (its 90% interval excludes zero, its 95% interval includes it). The remaining channel parameters include zero at both levels. The paper adopts the 90% interval as its inference standard. Panel B compares the country pairs bootstrap with the wild cluster bootstrap for the cumulative TFP statistic. With only five contributing DEM → AUT episodes the pairs bootstrap over-rejects, while the wild bootstrap, which does not degenerate with few clusters (MacKinnon and Webb, 2017), gives $p = 0.117$ for the DA cumulative effect. The cumulative five-year response is not distinguishable from zero. The inferential anchor is the impact-year contrast ($\hat{\beta}_0^{\text{DA}} = -0.042$, $t = -3.8$).

H Multi-period staggered estimators

The primary specification in Section 6.4 restricts attention to the first coup in each country and treats the country panel as truncated at the year of the second coup. This first-event design has two foundations. The first is the option-value framework of Section 3, which models the firm's adoption decision around the *first* major institutional shock under uncertainty. The reduced-form regression with a single shock indicator per country corresponds directly to that theoretical object. Estimating an average effect across multiple shocks per country would conflate the first-shock identification problem the theory speaks to with the marginal-shock identification problem it does not. The second foundation is econometric: with the Bjørnskov and Rode (2020) panel admitting up to 8 coup events per country over the 1960–2022 window, post-event horizons of the first coup overlap with pre-event horizons of subsequent coups in any country with two or more events, contaminating the estimated β_h in the absence of a staggered-event correction.

This appendix complements the primary first-event specification with two multi-period staggered estimators that admit subsequent coup events as additional treatments. The Callaway and Sant'Anna (2021) estimator assumes treatment is absorbing (each unit can be treated once). The de Chaisemartin and D'Haultfœ uille (2024) estimator allows treatment to switch on and off, the empirically appropriate specification for democracy–autocracy cycling. Both are reported as robustness for the primary first-event design, not as replacements, on theoretical grounds (option-value first-shock correspondence) and on empirical grounds documented in Subsections H.1–H.2 below.

H.1 Cycling diagnostics

We define the multi-period treatment indicator D_{it} as the cumulative-DA-event flag: $D_{it} = 1$ if country i has experienced at least one DEM $\rightarrow$ AUT coup on or before year t *and* is currently in an autocracy (regime$_{it} \in \{3, 4, 5\}$ in the Bjørnskov–Rode coding), and zero otherwise. The indicator switches off when a country returns to democracy after a previous DA event and switches back on at a subsequent DA event. Table 18 summarizes the resulting switch counts across the panel.

Table 18: Cycling diagnostics for the multi-period DA treatment indicator

Quantity	Count
Total switch-on events ($0 \to 1$)	65
Total switch-off events ($1 \to 0$)	47
Countries that switch on at least once	40
Countries that never switch	164
Total country-years in treated state ($D_{it} = 1$)	776
Total country-years in control state ($D_{it} = 0$)	∼10,300

Notes. 40 of 204 countries in the BR panel experience at least one DEM → AUT event over 1960–2022. Across those 40 countries, the panel records 65 switch-on events. The difference of $65 - 40 = 25$ reflects countries that cycle (return to democracy and re-enter autocracy through a subsequent DA event). 47 switch-off events occur where a previously-treated country returns to democracy.

The cycling rate is substantial. Argentina, Nigeria, and Honduras each cycle three times within the 1960–2022 window. Bolivia, Ecuador, Ghana, Guatemala, Niger, and Pakistan each cycle twice. This pattern violates the absorbing-treatment assumption underlying the Callaway–Sant'Anna estimator and motivates the de Chaisemartin–D'Haultfœuille non-absorbing alternative as the correctly-specified comparison.

Table 19 compares the three estimators on the impact-year TFP effect for DEM → AUT events.

Table 19: Three-estimator comparison: impact-year TFP effect

Estimator	Treatment process	$\hat{\beta}$	SE	p-value
First-event LP (baseline, §6.4)	Single-event per country	−0.042	0.012	0.001
Callaway–Sant'Anna (`csdid`)	Absorbing staggered	−0.032	0.012	0.008
de Chaisemartin–D'Haultfœuille (`did_multiplegt_dyn`)	Non-absorbing switching	−0.015	0.010	0.135
		Avg. effect	SE	95% CI
Callaway–Sant'Anna post-treatment average	Absorbing	−0.050	0.028	[−0.105, +0.004]
dCDH avg. total effect across post-shock horizons	Non-absorbing	−0.039	0.025	[−0.087, +0.010]

Notes. The first-event LP coefficient is the DEM → AUT BBG-decomposition estimate at $h = 0$ reported in Section 6.4. Callaway–Sant'Anna uses the first DA year per country as the cohort variable and includes the lagged log of GDP per capita as the time-varying covariate. Controls are the 164 never-switched countries. dCDH uses the non-absorbing switching indicator D_{it} defined in Subsection H.1 with five post-treatment effects, three placebo periods, and country-cluster standard errors. Both staggered estimators are implemented with default Liu et al. (2024) cluster-robust asymptotic standard errors.

The three estimators converge in magnitude. The first-event LP returns −0.042, csdid returns −0.032 at the impact year (average post-treatment effect −0.050), and dCDH returns −0.015 at the impact year (average total effect −0.039). All point estimates lie in the range [−0.050, −0.015], with the dCDH impact-year estimate the smallest in magnitude but with the same sign as the baseline.

H.2 Identification–power trade-off

The Callaway–Sant'Anna estimator gains identification power by exploiting the full pool of 164 never-treated control countries against each first-DA cohort, but at the cost of treating DA as absorbing in a setting where Argentina, Nigeria, and Honduras each cycle three times. The de Chaisemartin–D'Haultfœuille estimator specifies the treatment process correctly, allowing D_{it} to switch on and off, but pays for that correctness in identification power. Of the 65 switch-on events in the panel, only 20 enter the dCDH estimand for the impact-year effect: the estimator requires that the post-shock treatment trajectory remain stable for the duration of the dynamic horizon (here, five years), and the high-cycling countries (ARG, NGA, HND, BOL, GHA, GTM, NER, PAK, MLI, ECU) are systematically discarded because their treatment status flips back to zero within the window. The control pool is unaffected (164 never-switched countries provide ample counterfactual), but the treated sample collapses from 40 first-DA-event countries to 20 stable-trajectory switchers.

This is an identification–power trade-off, not evidence that the underlying effect is null. The dCDH joint nullity test on the post-shock horizons returns $p = 0.74$, but the point-estimate magnitude (-0.039 averaged across horizons, with individual horizon estimates ranging from -0.015 to -0.042) is aligned with the first-event LP baseline of -0.042 and with the csdid post-treatment average of -0.050. The convergence of magnitude with dispersion of precision is the pattern expected under the trade-off: the correctly-specified estimator preserves the signed-magnitude signal but loses statistical power because the non-absorbing structure of the empirical treatment process reduces the effective treated sample.

H.3 dCDH sample composition

Table 20 lists the 22 switchers that satisfy both the post-shock trajectory stability requirement (treatment remains on through $t+5$) and the TFP coverage requirement (non-missing $\ln \text{TFP}$ at all of $t-1$ through $t+5$). The dCDH estimator uses 20 of these in the impact-year identification, the remaining two being lost on a control-availability check that the estimator applies internally.

Table 20: Composition of the dCDH contributing sample

Country	Switch year	DA events at switch	First DA event?	Canonical recovered
ARG	1966	2	No (subsequent)	**Onganía**
ARG	1976	3	No (subsequent)	**Videla (*Proceso*)**
BRA	1964	1	Yes	Castelo Branco
BDI	1997	1	No	—
BDI	2015	2	No (subsequent)	—
CAF	2003	1	Yes	—
CHL	1973	1	Yes	Pinochet
ECU	1964	1	No	—
FJI	2000	1	Yes	Speight
GRC	1967	1	Yes	Colonels
HND	1973	2	No	—
MRT	2009	1	No	—
NGA	1966	1	Yes	Aguiyi-Ironsi
NGA	1984	2	No	**Buhari (lag-1)**
NGA	2003	2	No	—
PER	1969	1	No	**Velasco (lag-1)**
PER	1990	1	No	Artificial switch (see notes)
SDN	1989	2	No	Bashir
THA	2014	3	No	**Prayut**
TUR	2016	2	No (subsequent)	—
URY	1973	1	Yes	Bordaberry
ZMB	1997	1	Yes	—

Notes. The 22 country-year switchers satisfy dCDH `effects(5)`: $D_{i,t-1}$=0, D_{it}=1, $D_{it'}$=1 for $t' \in (t, t+5]$, plus non-missing $\ln$ TFP at $t-1$ through $t+5$. Two are lost in the internal control-availability check (dCDH reports 20). Eight overlap with the baseline first-event sample, the remaining fourteen are subsequent DA events that the first-event design omits. Five canonical post-1960 democracy-destroying coups appear as subsequent-event switchers: Onganía 1966 (Argentina), Videla 1976 (Argentina), Velasco 1969 (Peru, lag-1 relative to the 1968 event), Buhari 1984 (Nigeria, lag-1 relative to 1983), and Prayut 2014 (Thailand). Other canonical events (García Meza 1980, Argentina 1961, Honduras 1963, Guatemala 1963/1982, Ghana 1972/1982, Pakistan 1977/1999, Mali 2012/2021) are excluded by the cycling diagnostic of §H.1: each is followed by a return to democracy within five years. PER 1990 enters via a BR regime-code transition unaccompanied by a coup event. The Fujimori 1992 *autogolpe* is classified as intra-autocracy under the same coding (footnote 7 in §6.4).

The dCDH joint nullity test on the five post-shock event-study coefficients returns $p = 0.74$, which under a default reading would be interpreted as failure to reject the absence of a treatment effect. That reading is incorrect here. The joint test allocates power across five horizons identified from 20 effective switchers, with non-zero point estimates at every horizon (ranging from -0.015 to -0.042) and a clustered standard error of roughly 0.01–0.03

per horizon, the test simply cannot distinguish a joint -0.030 alternative from zero with the available identification mass. The cumulative-total-effect p-value on the -0.039 average over $[-0.087, +0.010]$ is the same observation in a single-parameter form.

The structural reason for the precision loss is documented in Subsections H.1 and H.2. The convergence of point- estimate magnitudes across the three estimators (first-event LP -0.042, csdid -0.050, dCDH -0.039 on the post-treatment average) is the inferentially relevant signal. The dispersion of p-values reflects the trade-off between treatment-process specification and identification power, not the presence or absence of a real effect.

The three estimators converge on the magnitude range $[-0.050, -0.039]$ for the average effect of DEM $\rightarrow$ AUT coups on impact-year and post-shock TFP. No estimator provides evidence that the underlying effect is null. The first-event LP design retained as the primary specification in the main text is justified on the theoretical correspondence with the option-value framework's focus on the first major institutional shock (Section 3) and on the empirical observation that the non-absorbing nature of the DA treatment process leaves the correctly-specified multi-period estimator under-powered relative to the first-event LP. The Callaway–Sant'Anna comparison provides the more powerful but mis-specified absorbing-treatment benchmark that returns a magnitude consistent with the first-event LP, and the de Chaisemartin–D'Haultfœuille comparison provides the correctly-specified but less powerful non-absorbing benchmark that returns the same signed magnitude with a wider confidence band.