Who Captures Trade-Network Growth? Institutional Capacity and the Absorption of Foreign Shocks

Jesse Schmolze*

May 8, 2026

Abstract

This paper investigates the heterogeneity of growth-shock transmission across bilateral trade networks. Standard models assume a uniform absorption rate for these spillovers. Instead, I hypothesize that institutional quality acts as a filter for shock propagation. Using a panel of 111 countries from 1996 to 2023, I apply Double Machine Learning to estimate the absorption rate while partialing out high-dimensional confounders, with WGI Political Stability as the institutional moderator. The DML estimates imply that a one standard deviation increase in institutional quality raises the spillover elasticity by 0.30, with the elasticity rising from approximately zero in the lowest stability tertile to 0.86 in the highest. Leave-one-out and CTOT IV specifications confirm the gradient and place the causal slope in the range [0.69, 1.07], indicating that the DML estimate is a conservative lower bound. The benefits of trade-led growth are conditional on institutional capacity.

*I disclose that Anthropic's LLM Claude assisted me in this project with the following tasks: data cleaning, coding the models, LaTeX troubleshooting, grammar checking, general writing advice and making figures. If you have any questions or concerns regarding my use of AI with this paper, I encourage the reader to reach out to jschmolze@wisc.edu.

Introduction

Consider two countries on opposite sides of the institutional distribution. Both participate in the same global trade network and their shared trading partners get an exogenous growth shock. The two countries face identical external conditions, yet one country converts their neighbors' shock into domestic growth while the other does not. The standard trade-spillover specification has no language for this asymmetry. It assumes that the absorption rate of partner growth is uniform across countries, with a single coefficient on the trade-weighted partner growth term. This paper challenges that assumption.

The literature on trade and growth, going back to Frankel and Romer Frankel and Romer (1999) and operationalized in Arora and Vamvakidis Arora and Vamvakidis (2005), treats absorption as a fixed parameter. Domestic growth is regressed on a trade-weighted partner growth term, a single coefficient is reported, and the implicit assumption is that all countries absorb partner growth at the same rate conditional on observables. The institutions-and-growth literature, established by Acemoglu, Johnson, and Robinson Acemoglu et al. (2001) and refined by Rodrik, Subramanian, and Trebbi Rodrik et al. (2002), has independently shown that institutional quality drives long-run development. These two literatures have rarely been brought together at the level of the absorption coefficient itself. The closest precedent is Coe, Helpman, and Hoffmaister Coe et al. (2009), who document that institutional quality moderates international R&D spillovers, but they study knowledge transmission rather than growth, and impose a potentially restrictive linear interaction. The question of whether the homogeneous-absorption assumption in the broader trade-growth literature is empirically defensible remains open.

This paper estimates a heterogeneous absorption function $\theta(D)$ across a panel of 111 countries from 1996 to 2023. The function maps own-country political stability to the rate at which partner growth shocks are absorbed into domestic growth. To handle the high-dimensional and highly collinear nature of institutional indicators, I implement the Double Machine Learning framework of Chernozhukov et al. (2018), extended to panel settings by Clarke and Polselli Clarke and Polselli (2026) and to heterogeneous treatment effects by Semenova and Chernozhukov Semenova and Chernozhukov (2017). To establish causality, I introduce two instruments. The first is a Leave-i-Out network instrument that severs the reflection feedback loop inherent to network spillover models. The second is a partner-weighted Commodity Terms of Trade instrument that exploits exogenous price shocks to partners' export baskets. The two instruments exploit fundamentally different sources of variation, and the Hansen J test fails to reject the null that both yield consistent estimates of $\theta(D)$.

The absorption gradient is positive, monotone, and economically large. Under the joint IV specification, $\theta(D)$ rises from approximately zero in the bottom tertile of the political stability distribution to above one in the top tertile. Countries below the Low-stability threshold are statistically disconnected from the global growth network, with absorption indistinguishable from zero. Countries above the High-stability threshold absorb more than one-for-one of partner growth shocks. This is consistent with a financial-accelerator mechanism in which firms borrow against expected revenue and expand capacity beyond the proportional response. The institutional gradient survives across all alternative governance indicators tested and across all alternative first-stage learners. The Anderson-Rubin weak-IV-robust confidence interval for the gradient lies entirely above the DML baseline estimate, indicating that the baseline was a conservative lower bound on the true causal effect.

This paper makes one main contribution: It documents that the homogeneous-absorption assumption in the trade-spillover literature is misspecified. Single-coefficient specifications estimate a sample-weighted average of a function that varies significantly across the institutional distribution, and policy advice derived from these models should be conditioned on institutional capacity rather than treating absorption as a uniform structural parameter.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature. Section 3 describes the data and the construction of the network spillover term. Section 4 outlines the empirical strategy. Section 5 presents the main results. Section 6 discusses the implications. Section 7 reports robustness checks. Section 8 concludes.

Literature Review

I begin the Literature Review by laying the foundational premises my paper is built upon. First, Acemoglu et al. (2001) is the canonical piece that argues for institutions driving economic growth. They reinforce this with another paper published a year later Acemoglu et al. (2002), that suggest the regions that transitioned to poor from rich between the 1500s and modern day have institutional problems as the primary drivers of that change. Taken together, these papers provide robust evidence that institutional quality is a primary determinant in macroeconomic growth. I build on them by narrowing the scope of my work specifically to how growth propagates from trade networks, and I highlight a more specific mechanism for which the results they describe may operate.

Hall and Jones (1999) further reinforce my approach, when they find "social infrastructure" is a deep determinant of variations of countries' productivity differences. Furthermore, Rodrik et al. (2002) find that institution quality is much more important than geography and trade for economic growth once you instrument properly. In my paper, I argue for a more nuanced approach than what Rodrik et al. took. I Instead suggest that the impact of the trade on economic growth is directly mediated by political stability, which suggests that treating them as separate forces doesn't provide a full outlook of the scenario.

In the literature, there has been some push back about the idea that institutions cause growth. Notably, Glaeser et al. (2004) argue that human capital, not institutions, is what colonizers brought over, which contradicts the results from the aforementioned Acemoglu papers. They claim that institutional measures are too "noisy" to separate from policy choices. While I do not aim to adjudicate this debate, my own results do not indicate that the institutional measures used here suffer from the noise Glaeser et al. emphasize. Still, the reader should recognize that although the foundations of this paper rest on a well-established literature, that literature is not without credible challenges.

In my paper, I choose to use Political Stability as the primary moderator of D. This is well supported by empirical literature. For example, Alesina et al. (1996) find instability significantly reduces growth in a cross country panel. Barro (1991) also includes political stability proxies in the widely regarded Barro growth regressions. Finally, Aisen and Veiga (2011) uses system GMM on a panel of 169 countries and confirms the negative effect of instability and explores the channels through which it operates. I build on their work by specifically analyzing how political stability affects how trade related growth propagates across borders.

As we pivot towards trade specifically, I want to introduce Acemoglu et al. (2003), which shows that weak institutions cause macroeconomic volatility, which notably includes worse response to shocks. This paper establishes a key notion of my work: bad institutions respond worse to growth shocks. Rodrik (1999) further supports this by analyzing how weaker institutions responded worse to the adverse external shocks of the 1970s. He argues that this fact explains a large portion of the post-1975 collapse in much of the developing world. These papers give strong support that institutions mediate how a country is affected by economic events exogenous to their economy.

Moving to the trade, Frankel and Romer (1999) discuss the flagship example of how trade openness has a positive causal effect on income. Feyrer (2019) updates the Frankel-Romer paper using time-varying air/sea transport costs, which confirms the elasticity holds when exploiting the within-country variation. Wacziarg and Welch (2008) also support this claim by arguing that liberalization episodes are followed by growth accelerations. These papers serve as the bedrock of my argument, but as you will see in my paper, the approach that trade causes income is more nuanced than a one-line cause and effect.

The mechanism of trade causing growth is laid out by Coe and Helpman (1995), where they construct foreign R&D capital using trade-weighted averages. They find that R&D capital leads to positive spillovers, and I also use their trade weighted methodology later in the paper (see methodology section). Keller (2004) finds that trade is the dominant channel through which growth-relevant knowledge moves across borders, suggesting there's a network effect other than purely trade in the mechanism. Arora and Vamvakidis (2005) do similar work to me when they regress trade-weighted partner growth and find a large coefficient. This is another direct precedent to my work, which involves them finding the average effect of my mechanism. I show that significant heterogeneity exists in the effect conditional on the quality of institutions.

In terms of the shock progression, Acemoglu et al. (2016) lay out the theoretical foundation for why network structure matters. Their original framework is about input-output networks within an economy, but you can broaden the scope of their work to the world economy and the same results hold true. This work provides a concrete theoretical basis on top of which I lay out my findings. Di Giovanni and Levchenko (2010) show trade linkages drive cross-country correlation in business cycles, which illustrates how my proposed hypothesis becomes quantitatively important. Finally, Caselli et al. (2020) argue that openness to trade mediates domestic GDP volatility, thus causing less uncertainty and more long run growth.

Moving onto Organizational Economics, it's important to define what "absorptive capacity" means. The origin of it comes from Cohen and Levinthal (1990), who argue that firms need internal capabilities to absorb external knowledge. In a macroeconomic environment, Abramovitz (1986) introduces "social capability" as the precondition for technological catch up. He argues that countries don't automatically converge to the frontier, they need the institutional and human capital infrastructure to absorb frontier knowledge.

This is formalized in the Foreign Direct Investment Literature. The most important work is Borensztein et al. (1998), which finds that FDI raises growth only in countries with sufficient human capital. This serves as an effect analogous to the one investigated in this paper. Alfaro et al. (2004) show that financial development is the key absorption channel for FDI, and they also demonstrate that this mechanism is multidimensional which informs my methodology throughout this paper. Durham (2004) generalizes Borensztein across multiple capital flows and absorption channels.

My closest precedent in terms of the trade network spillovers comes from R&D literature. Coe and Helpman (1995) was already cited as my foundational basis, but they estimate a uniform spillover coefficient with no institutional moderation. I argue that that notion is restrictive and the effect varies significantly based on institutional quality. They subsequently examined this exact mechanism across R&D channels, where they confirm that institutional differences are vital determinants of TFP Coe et al. (2009). They specifically find that countries with the most robust institutional frameworks extract significantly larger productivity gains from international knowledge transmission. The existence of this effect in R&D highly motivates my work to look at it through a lens of international growth. I add to their work by doing three things: 1) Looking at growth spillovers and not R&D spillovers; 2) They use linear restrictions to the model, I use a more flexible approach that accounts for higher dimensionality in the data. In the end, I find there are no higher dimensional effects which validates the linear restriction; 3) I address the reflection problem in network spillovers by using a network leave-one-out instrument.

I finish up the literature review by motivating my methodological choices. The most important decision I use is using Double Machine Learning (DML) Chernozhukov et al. (2018) as my main method of estimation. As noted by Belloni et al. (2014), naive applications of machine learning to causal estimation suffer from regularization bias, where the shrinkage of coefficients toward zero introduces significant bias into the structural parameter. It introduces Neyman-orthogonal moment conditions and cross-fitting, which when combined, allows me to use ML predictors for nuisance functions while preserving $\sqrt{n}$ inference on the structural parameters. DML's two innovations eliminate the regularization bias asymptotically. This is what lets me flexibly control for high dimensional institutional confounders without imposing a parametric form.

I also use the Clarke and Polselli (2026) recently developed methods for implementing DML for static Panel models with fixed effects. This is the direct methodological background of my work. It extends DML to two-way fixed effects panels, which is the exact setting of my paper. Finally, I use the methods laid out in Chernozhukov et al. (2018) to determine the heterogeneous DML treatment effects, which is what I use to estimate $\theta(D)$. Without this work, I can only estimate the average effect and my analysis fails.

Moving to the IVs I use, I would like to motivate 2 specific IVs. The first is the LOO instrument. This seeks to sever an endogenous feedback loop in my analysis. Manski (1993) finds that in a linear-in-means peer effects model, you cannot separately identify endogenous effects from contextual effects and correlated effects. This is the statement that my identification problem my LOO addresses. This is addressed by Bramoullé et al. (2009). They show that you can identify endogenous peer effects using the characteristics of peers' peers as instruments. By excluding country i from the spillover calculation of country j, I can exploit network non transitivity to construct an instrument that varies with j's environment but is orthogonal to i's outcome.

I additionally use commodity price shocks as exogenous variation. The strategy was first formulated by Deaton and Miller (1995), where they analyzed the macroeconomies of Sub-Saharan Africa. The methodology of the index I'm using is laid out in Kebhaj (2020). The relevance condition is supported by Fernández et al. (2017), where they explain that global commodity shocks explain a large share of business cycle variation in emerging markets. Drechsel and Tenreyro (2018) back this up by showing that commodity price shocks transmit strongly to growth in commodity exporting countries.

Data

In this section, I will go over the data I use in my analysis. I will cover the panel structure of my dataset, my four sources, and the two engineered network variables I use. After parsing through all the data and restrictions, I end up with $N = 2{,}103$ country-years, 111 countries, from 1996 to 2023.

3.1 Sample and Panel Structure

My panel samples 120 panel countries from the PWT. I restrict this panel to 111 with complete WGI coverage post 1996. I use annual frequency throughout all my variables. The 1996 start is a binding constraint imposed by WGI availability. The Penn World Table data goes back to 1970, but the moderator D does not. My dependent variable is $g_{it} = \%$ change in real GDP per capita from PWT. Spillover terms use $t-1$ partner growth and $t-1$ trade weights to address simultaneity at the within-year level.

3.2 Macroeconomic Indicators

I use PWT 10.01 Feenstra et al. (2015) for real GDP per capita, population, investment share, and government consumption share. I winsorize investment share and government consumption share at the 1/99% to limit the influence of outliers driven by hyperinflation episodes and reporting anomalies. I also add a standard convergence control $\ln y_{i,t-1}$.

3.3 Bilateral Trade Network

My trade data is sourced from the IMF Direction of Trade Statistics (DOTS) International Monetary Fund (2024); Marini et al. (2018), which provides bilateral nominal exports and imports for the country pairs in my sample. I define the trade weighted variable

$$W_{ijt} = \frac{X_{ijt} + M_{ijt}}{\sum_{k \neq i}(X_{ikt} + M_{ikt})},$$

where X_{ijt} is exports from i to j and M_{ijt} is imports of i from j. So W_{ijt} is the share of j in i's total bilateral trade in year t. By construction, $\sum_{j\neq i} W_{ijt} = 1$. I use $W_{ij,t-1}$ for the spillover construction because it reduces simultaneity. I would like to further acknowledge that $W_{ijt} \neq W_{jit}$ because the denominator is country specific.

3.4 Institutional Quality

I use WGI's indicators Kaufmann et al. (2010). WGI provides six institutional dimensions, all measured as standardized z-scores with mean zero and unit variance. My moderator choice is D_{it} = WGI Political Stability. See Section 4.1 where I justify this choice with my pre-screening exercise. I had an imputation issue prior to 2002, so I linearly interpolated the off years.

3.5 Engineered Network Variables

In this section, I define two key variables. The first is the network spillover term

$$\tilde{g}_{it} = \sum_{j\neq i} W_{ij,t-1} \cdot g_{j,t-1}.$$

This is the trade weighted average of partner growth rates lagged one period. It captures the growth shock transmitted to country i through its bilateral trade exposures. Partner institutional quality is defined in an analogous way:

$$\tilde{D}_{it} = \sum_{j\neq i} W_{ij,t-1} \cdot D_{jt},$$

where D_{jt} is country j's domestic political stability. This is the trade weighted average of partners' political stability. This captures the institutional environment of i's trade network. Note that when I use the LOO variant in Section 4, $\tilde{g}_{it}$ becomes $\tilde{g}^{-i}_{j,t-1}$. This excludes country i from j's spillover calculation. Finally, the partner weighted commodity export price shock is constructed analogously using IMF CTOT data Adler and Magud (2015).

3.6 Summary Statistics

Table 1: Summary statistics on the exact DML sample (N = 2,103, 111 countries, 1996–2023).

Variable	Mean	SD	Min	Max
g_{it} (% real GDP/cap growth)	3.05%	6.10%	−26.8	+45.1
$\tilde{g}_{it}$ (trade-weighted spillover)	2.93%	2.34%	−6.30	+18.98
D_{it} (WGI Political Stability, z-score)	0.16	0.93	−2.81	+1.76

The summary statistics confirm that the spillover term $\tilde{g}_{it}$ is much less volatile than own-country growth g_{it} (SD of 2.34% vs. 6.10%), which is expected since the trade-weighted average smooths idiosyncratic partner shocks. Political Stability is approximately mean-zero by construction, with substantial cross-country dispersion.

Together, these data permit estimation of how growth shocks propagate through bilateral trade networks, conditional on the institutional quality of the receiving country. The next section formalizes the empirical strategy.

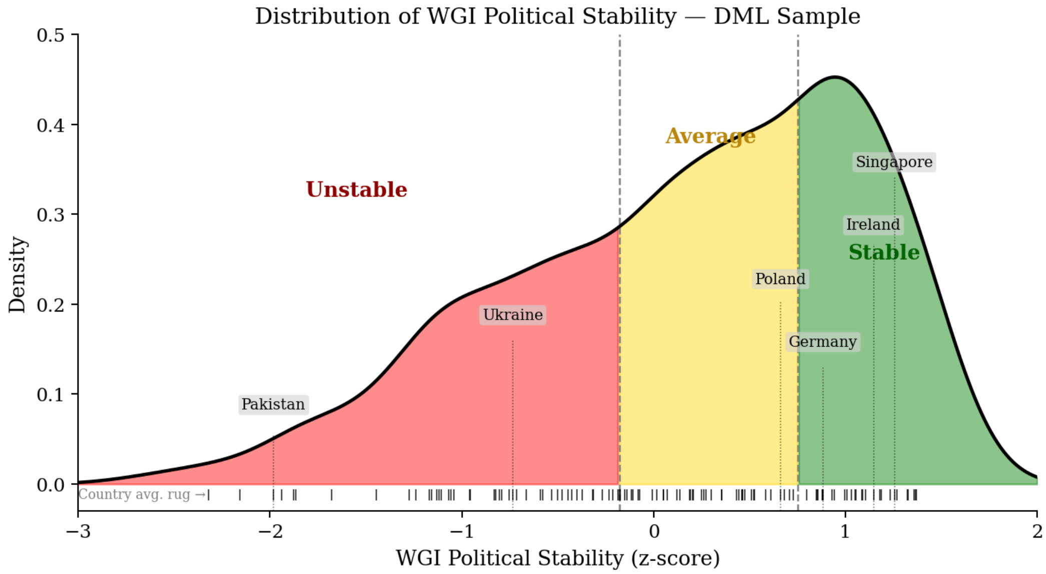

Figure 1: **Distribution of WGI Political Stability and Tertile Classification.** Kernel density estimate of WGI Political Stability z-scores across the full DML estimation sample ($N = 2{,}103$ country-years). The distribution is segmented into three regimes based on sample tertiles: Unstable (red, $D < -0.18$), Average (yellow, $-0.18 \leq D \leq 0.75$), and Stable (green, $D > 0.75$). The rug plot along the x-axis represents the 111 unique country averages, with selected country markers (Pakistan, Ukraine, Poland, Germany, Ireland, Singapore) provided for geographic context.

Empirical Strategy

4.1 Choosing the Moderator

Quantifying institutional quality is an inherently difficult task. The major governance indices are highly intercorrelated, with cross-correlations exceeding 0.7 across the WGI cluster and reaching 0.95 between WGI Voice and Accountability and the V-Dem democracy indices. Selecting a single dimension as the moderator therefore requires an empirical criterion.

I select D_{it} by running the two-way fixed-effects regression across each of eight candidate institutional indicators. The candidates include all six WGI dimensions and the two principal V-Dem democracy indices. Results are summarized in Table 2.

The estimates reflect significant variability in both the economic and statistical significance of these indicators. Notably, WGI Political Stability and WGI Regulatory Quality stand out as having large economic and statistical significance ($\hat{\beta} = 1.791$, $p < 0.01$ and $\hat{\beta} = 2.535$, $p < 0.05$ respectively). This aligns with intuition, as one can imagine how a smooth

Table 2: Screening Candidate D Variables: Two-Way FE Regressions of Growth on Institutional Indicators

D variable	β	SE	p-value	Sig.
WGI Political Stability	**1.791**	**0.566**	**0.002**	***
WGI Control of Corruption	1.439	1.202	0.231	
WGI Gov. Effectiveness	1.183	1.094	0.280	
WGI Regulatory Quality	2.535	1.042	0.015	**
WGI Rule of Law	1.513	1.094	0.167	
WGI Voice & Accountability	1.104	0.925	0.232	
V-Dem Liberal Democracy	0.733	1.313	0.577	
V-Dem Electoral Democracy	−0.021	1.186	0.986	
Joint specification (Spec 9):				
WGI PV	**1.833**	**0.577**	**0.001**	***
V-Dem Liberal Democracy	−1.669	2.275	0.463	

Notes: Each row reports the coefficient on D_{it} from a separate two-way fixed-effects regression of GDP per capita growth on D_{it} and the control set X_{it}: ln GDP/cap (lag), population growth, human capital, investment share, government consumption share, and trade openness. Country and year fixed effects included. Standard errors clustered by country. WGI sample: $N = 3{,}258$, 120 countries, 1996–2023. V-Dem sample: $N \approx 5{,}540$, 118 countries, 1970–2023. R^2: 0.236–0.241 (WGI specs); 0.169–0.170 (V-Dem specs).
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

political environment and strong regulatory frameworks act as primary amplifiers for positive foreign shocks. Conversely, measures of democratic participation, such as V-Dem's Electoral Democracy Coppedge et al. (2024); Pemstein et al. (2025), yield statistically insignificant coefficients near zero. This implies that the mechanics of elections matter far less for shock absorption than the stability of the state itself.

I select WGI Political Stability as the primary moderator D_{it} for two reasons. First, it is the only indicator significant at the 1% level, and its coefficient remains stable when jointly specified with V-Dem Liberal Democracy ($\hat{\beta} = 1.833$, $p = 0.001$). Second, it has the lowest cross-correlation with the V-Dem indices ($r = 0.53$ to 0.61), making it the most conceptually distinct from electoral participation channels and the cleanest available measure of the institutional capacity dimension this model is designed to identify. Section 7 confirms that the moderating gradient documented below is not specific to this choice. The result survives across all six WGI dimensions.

4.2 Double Machine Learning

Moving on to the DML, I implement a partially linear model with a heterogeneous treatment effect on the network spillover term:

$$g_{it} = \theta(D_{it}) \cdot \tilde{g}_{it} + f(X_{it}) + \alpha_i + \lambda_t + \epsilon_{it} \tag{1}$$

I parameterize $\theta(D)$ as linear in the moderator:

$$\theta(D_{it}) = t_0 + t_1 D_{it} \tag{2}$$

The linear specification is not a simplification. I test quadratic and cubic augmentations in Section 7 and fail to reject linearity in both cases ($p = 0.871$ and $p = 0.333$ respectively). The implied vertex of the quadratic specification falls at $D^* = -9.1$, far outside the data range, confirming that $\theta(D)$ is monotone over all observed institutional levels.

I follow Chernozhukov et al. (2018) for the foundational DML framework. Their work establishes Neyman-orthogonal moment conditions paired with cross-fitting, which together let me use ML predictors for the nuisance functions while preserving valid inference on the structural parameter. Without Neyman orthogonality, regularization bias from the ML first stage would contaminate the second-stage estimate of θ.

I also use Clarke and Polselli (2026) recent methods for implementing DML for static panel models with fixed effects. This is the direct methodological background of my work. It extends DML to two-way fixed effects panels, which is the exact setting of my paper. Finally, I use the methods laid out in Semenova and Chernozhukov (2017) to determine the heterogeneous DML treatment effects, which is what I use to estimate $\theta(D)$. Without this work, I can only estimate the average effect and my analysis fails.

In practice, I implement 5-fold cross-fitting with Ridge nuisance learners, partitioning the sample by country to preserve panel structure. The nuisance functions are:

$$\hat{V}_{it} = \tilde{g}_{it} - \hat{m}^{(-k(it))}(X_{it}), \qquad \hat{U}_{it} = g_{it} - \hat{\ell}^{(-k(it))}(X_{it}) \tag{3}$$

where $k(it)$ denotes the fold containing observation it, and $\hat{m}$ and $\hat{\ell}$ are estimated on all folds excluding $k(it)$. The second-stage moment condition is:

$$\mathbb{E}\left[\left(\hat{U}_{it} - \theta(D_{it})\hat{V}_{it}\right) \cdot \hat{V}_{it} \cdot (1, D_{it})'\right] = 0 \tag{4}$$

A natural concern with this framework is that D_{it} may proxy for recent growth performance. If politically stable countries are simply on faster growth trajectories, the estimated gradient t_1 would be misleading. I test this directly by regressing D on lagged growth after two-way FE removal. The resulting $R^2 = 0.0015$, with raw correlations near zero. Within-country variation in political stability is near-orthogonal to lagged growth, ruling out the proxy story.

A second concern is that the DML machinery may itself be driving the result. To address this, I compare cross-fitted Ridge residuals against classical Frisch-Waugh-Lovell OLS residuals. The two estimators agree to within 1.3% on t_1, and the underlying residual series correlate at $r = 0.999$. In this setting, the value of DML is robust inference under potential nonlinearity in $f(X)$, not a different point estimate. The result survives whether or not the reader accepts ML.

4.3 Causal Identification

The DML framework above removes bias from high-dimensional observables, but it does not address the endogeneity of the network spillover term $\tilde{g}_{it}$. Even after partialing out X and the two-way fixed effects, three distinct threats to causal identification remain.

The first is common factor bias. Global and regional shocks such as financial crises and commodity cycles hit g_{it} and $g_{j,t}$ simultaneously. Year fixed effects absorb the global mean, but regional or bilateral correlated errors persist.

The second is the reflection problem. Country i's growth affects partner j's growth through j's trade with i, which then feeds back into $\tilde{g}_{it}$ through the bilateral weight matrix. Lagging the trade weights to $W_{ij,t-1}$ breaks one direction of this loop, but it does not break the circular network structure underlying it.

The third is weight endogeneity. Countries form trade relationships with fast-growing partners, so $W_{ij,t-1}$ reflects prior joint growth dynamics and is not randomly assigned.

To address these threats, I construct two instruments for $\tilde{g}_{it}$. The first is a leave-i-out network instrument that severs the reflection feedback loop. The second is a partner-weighted commodity terms of trade instrument that exploits exogenous price variation in partners' export baskets. The instruments are constructed independently and are nearly orthogonal in the data ($\text{corr}(Z^{LOO}, Z^{CTOT}) = -0.014$), which lets me jointly estimate the system and test the over-identifying restriction. I detail each instrument in turn.

4.4 Leave-i-Out Network Instrument

I construct my first instrument Z_{it}^{LOO} as the trade-weighted average of partner growth, where each partner's spillover is computed with country i excluded from the calculation:

$$Z_{it}^{LOO} = \sum_{j \neq i} W_{ij,t-1} \cdot \tilde{g}_{j,t-1}^{-i} \tag{5}$$

where $\tilde{g}_{j,t-1}^{-i}$ is partner j's trade-weighted growth after removing i's contribution to j's network:

$$\tilde{g}_{j,t-1}^{-i} = \frac{\tilde{g}_{j,t-1} - W_{ji,t-1} \cdot g_{i,t-1}}{1 - W_{ji,t-1}} \tag{6}$$

This severs the feedback loop where country i's growth affects partner j, which then reflects back to i through the spillover term. The construction follows the network leave-one-out propagation design of Acemoglu et al. (2016), and the underlying identification logic is that of Bramoullé et al. (2009), who show that endogenous peer effects can be identified through the non-transitive structure of the network. By excluding i from j's spillover calculation, Z_{it}^{LOO} varies with j's exogenous environment but is orthogonal to i's contemporaneous outcome.

The instrument is strong. The first-stage cluster-robust Kleibergen-Paap F statistic is 51.95, well above the Stock-Yogo 10% critical value of 16.38. Throughout the IV results, I report Anderson-Rubin weak-IV-robust confidence intervals to ensure inference does not depend on first-stage strength.

4.5 Commodity Terms of Trade Instrument

I also explore using a second instrument that exploits exogenous price shocks to partners' export baskets as an alternative source of causal variation. I construct the partner-weighted commodity terms of trade instrument, Z_{it}^{CTOT}, as:

$$Z_{it}^{CTOT} = \sum_{j \neq i} W_{ij,t-1} \cdot \mathrm{CTOT}_{j,t}^{\mathrm{export}} \tag{7}$$

This strategy was first formulated by Deaton and Miller (1995), where they analyzed the macroeconomies of Sub-Saharan Africa. The methodology of the index I'm using is laid out in Kebhaj (2020). The relevance condition is supported by Fernández et al. (2017), where they show that global commodity shocks explain a large share of business cycle variation in emerging markets. Drechsel and Tenreyro (2018) back this up by showing that commodity price shocks transmit strongly to growth in commodity exporting countries.

The mechanism is straightforward. A commodity price boom in partner j raises j's growth, which then transmits to i through the trade-weighted spillover term. The exclusion restriction holds for two reasons. First, year fixed effects absorb the global commodity super-cycle, which accounts for 13.9% of total CTOT variance in my sample. The remaining 49.1% of variance is country-specific, and this is what the IV identifies from. Second, I explicitly avoid using a country's own CTOT as an instrument, because it violates the exclusion restriction by directly affecting domestic income.

4.6 Joint Estimation and Inference

I implement DML-IV by adding a third nuisance function for each instrument, residualizing it against the high-dimensional controls in the same cross-fitted manner as $\tilde{g}$ and g:

$$\hat{\psi}_{it}^{LOO} = Z_{it}^{LOO} - \hat{r}^{LOO,(-k(it))}(X_{it}), \qquad \hat{\psi}_{it}^{CTOT} = Z_{it}^{CTOT} - \hat{r}^{CTOT,(-k(it))}(X_{it}) \tag{8}$$

The second-stage moment condition replaces the OLS moment with an IV moment for each instrument:

$$\mathbb{E}\left[\left(\hat{U}_{it} - \theta(D_{it})\hat{V}_{it}\right) \cdot \hat{\psi}_{it}^{m} \cdot (1, D_{it})'\right] = 0, \quad m \in \{LOO, CTOT\} \tag{9}$$

I report results from each instrument individually as well as a joint estimator that stacks both moments. With two instruments and one endogenous regressor per moment, the system is over-identified, which lets me test instrument consistency through the Hansen J statistic:

$$J = n\,\bar{g}(\hat{\theta})'\,\hat{\Omega}^{-1}\,\bar{g}(\hat{\theta}) \;\sim\; \chi^2_{q-p} \tag{10}$$

The joint estimator yields $J = 1.384$ with $p = 0.501$, indicating the two instruments produce statistically consistent estimates of $\theta(D)$ and providing direct empirical support for the validity of both identification strategies.

I additionally report a lag-structure placebo for each instrument in Section 7, regressing the instrument on lagged own growth after fixed-effect removal to test for backdoor channels into the outcome.

For inference, I cluster standard errors at the country level throughout following Colin Cameron and Miller (2015). For the IV specifications, I additionally report Anderson-Rubin weak-IV-robust confidence intervals to ensure that inference on t_1 does not depend on first-stage strength. As a final robustness check, I also confirm that the second-stage estimates are stable across alternative first-stage learners (OLS, Ridge, Lasso, and Random Forest), with the implied range for t_1 documented in Section 7.

Results

5.1 Baseline DML

I begin with the DML baseline, which establishes the institutional moderating gradient before introducing causal identification. The full sample is 2,103 country-year observations across 111 countries from 1996 to 2023.

Table 3: Linear Specification $\theta(D) = t_0 + t_1 D$

Parameter	Estimate	SE	p-value
Average effect θ_{avg}	0.444	0.128	< 0.001
Intercept t_0	0.427	0.117	< 0.001
Slope t_1	0.295	0.111	0.008

The slope t_1 is positive and significant at the 1% level, confirming the central hypothesis that absorption of partner growth is not uniform but is moderated by institutional quality. Each one standard deviation increase in WGI Political Stability adds roughly 0.30 percentage points to the spillover elasticity.

To make the magnitudes concrete, I evaluate $\theta(D)$ at the center of each stability tertile. The results are shown in Table 4.

Table 4: Spillover Elasticity $\theta(D)$ by Stability Tertile

Tertile	Range	$\theta(D)$	Significance
Low	$D < -0.18$	0.08	n.s. (indistinguishable from 0)
Mid	$-0.18 \leq D \leq 0.75$	0.59	*** ($p < 0.01$)
High	$D > 0.75$	0.86	*** ($p < 0.01$)

The economic interpretation is direct. Stable countries absorb roughly 86% of partner growth shocks one-for-one. Countries in the bottom tertile of the stability distribution are effectively disconnected from the global growth network, with θ statistically zero. The threshold structure aligns with intuition: a country needs a baseline level of institutional capacity before it can convert external opportunities into domestic growth, and below that threshold the spillover channel does not operate.

I emphasize that this is a baseline. The DML framework removes bias from observable confounders but does not address the endogeneity of the spillover term itself. The instrumental variable estimates that follow refine these results under causal identification. The full results are summarized in Figure 2

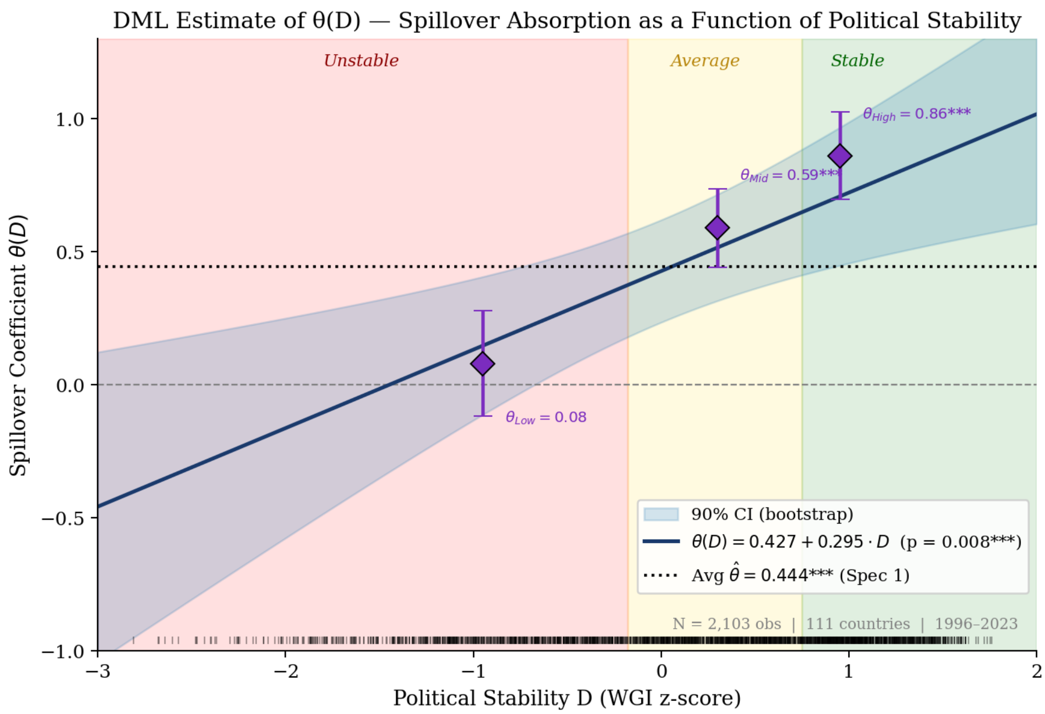

Figure 2: **DML Estimate of Spillover Absorption as a Function of Political Stability.** The blue line represents the estimated linear function $\theta(D) = 0.427 + 0.295 \cdot D$, with a 90% confidence band shaded in light blue. Orange diamonds represent the tertile point estimates with 90% confidence intervals. Tertile cuts are at $D = -0.18$ (33rd percentile) and $D = 0.75$ (67th percentile). The horizontal dotted line at $\theta = 0.444$ marks the average treatment effect. The rug plot along the x-axis shows the distribution of observed political stability values across the sample. N = 2,103 observations across 111 countries from 1996 to 2023.

5.2 Leave-i-Out Network IV

The first-stage diagnostics for the LOO instrument are strong. The cluster-robust Kleibergen-Paap F statistic is 51.95, well above the Stock-Yogo 10% critical value of 16.38. This rules out weak-instrument concerns and supports the use of standard Wald inference, though I additionally report Anderson-Rubin confidence intervals for transparency.

Table 5: Second-Stage IV Results

Parameter	Estimate	SE	p-value
Intercept t_0	0.327	0.511	0.522
Slope t_1	0.880	0.644	0.172
Anderson–Rubin 95% CI for t_1: [0.694, 1.067]			

The point estimate of t_1 jumps from 0.295 in the DML baseline to 0.880 under causal identification, roughly a threefold increase. Crucially, the AR-robust confidence interval for t_1 lies entirely above the DML baseline estimate, indicating that the baseline was a conservative lower bound on the true causal gradient. The Wald p-value of 0.172 reflects the increased standard errors typical of IV estimation, but the AR-robust interval excludes zero by a wide margin and is the more reliable inference object here.

The direction of the bias is informative. If reflection were the dominant source of endogeneity, IV correction would attenuate the estimate, since reflection mechanically inflates the OLS-style coefficient through the feedback loop. Instead, IV steepens the gradient. The most plausible explanation is that the DML baseline absorbs part of the institutional gradient through the high-dimensional control set. Countries with weak institutions occasionally experience volatile growth driven by commodity booms or unsustainable borrowing, and the controls partial out this variation alongside the genuine spillover channel. Once the LOO instrument isolates purely exogenous variation in partner growth, the role of institutions in gating the channel becomes more transparent and the gradient sharpens.

Table 6: IV-Implied Spillover Elasticity $\theta(D)$ by Stability Tertile

Tertile	Range	$\theta(D)$
Low	$D < -0.18$	≈ 0 (robustly zero)
Mid	$D \approx 0.30$	0.59
High	$D > 0.75$	1.16

The High-tertile result is economically striking. Highly stable countries absorb more than one-for-one of partner growth. This suggests that strong institutions allow firms to leverage external demand into capacity expansion rather than simply passing it through. The Low-tertile result remains near zero under IV. This demonstrates how in the absence of basic political stability, the trade-network growth channel is effectively closed.

5.3 Commodity Terms of Trade IV

I now turn to the CTOT instrument. The first-stage diagnostics are weaker than those for LOO, with the partial F statistic at 6.69 by Wald. This is below the Stock-Yogo threshold, so I treat the CTOT-alone estimates as suggestive rather than primary.

Table 7: Second-Stage CTOT IV Results

Parameter	Estimate	SE	p-value
Intercept t_0	-0.510	0.978	0.602
Slope t_1	1.278	1.737	0.462

The CTOT estimate of t_1 is large and positive but imprecise. The point estimate of 1.278 is well above the DML baseline of 0.295 and above the LOO estimate of 0.880, but the wide standard errors prevent a clean significance test on its own.

Even with the imprecision, the CTOT result is informative for two reasons. First, the sign and rough magnitude align with the LOO estimate, which matters because the two instruments exploit fundamentally different sources of variation. LOO uses bilateral network structure; CTOT uses exogenous commodity price shocks. The fact that both point in the same direction makes it unlikely that the LOO result is an artifact of one specific identification strategy. Second, the agreement sets up the joint estimator in the next subsection, where the two instruments are stacked to recover precision while testing for over-identification.

I emphasize that I do not interpret the CTOT estimate in isolation. The instrument is most useful as a robustness check on the LOO design and as one of two moments in the joint specification.

5.4 Joint Estimation

The cleanest causal claim in the paper comes from the joint estimator, which stacks the LOO and CTOT moments. With two instruments and one endogenous regressor per moment, the system is over-identified, which lets me both gain precision and formally test instrument consistency.

Table 8: Joint LOO + CTOT IV Estimates

Parameter	Estimate	SE	p-value
Intercept t_0	0.067	0.529	0.899
Slope t_1	0.925	0.646	0.152
Hansen $J = 1.384$ ($df = 2$, $p = 0.501$)			

The Hansen J test fails to reject the null that both instruments are consistent with the same value of $\theta(D)$. This is direct empirical support for the validity of both identification strategies. If either instrument were violating exclusion in a way that biased θ, the two would disagree and J would reject.

The joint slope of 0.925 sits between the LOO estimate of 0.880 and the CTOT estimate of 1.278, as expected from a precision-weighted combination. The joint intercept collapses toward zero, which is consistent with the Low-tertile finding that countries below the stability threshold have no detectable absorption.

Table 9 summarizes the full set of estimators alongside the joint $\theta(D)$ evaluated at the tertile centers. Across every IV specification, the slope t_1 is substantially larger than the DML baseline estimate of 0.295. The honest range for the causal t_1 is $[0.694, 1.067]$, defined by the AR-robust CI from the LOO specification. This range is uniformly above the DML baseline, confirming that the DML estimate is a conservative lower bound on the true institutional gradient. Panel B evaluates the joint $\theta(D)$ at the tertile centers and gives the cleanest summary of the paper's main result.

Table 9: Summary of $\theta(D)$ Estimates Across Specifications

Specification	t_0 (SE)	t_1 (SE)	$p(t_1)$	Diagnostic
Panel A: Estimator comparison				
DML (linear)	0.427 (0.117)	0.295 (0.111)	0.008	
IV: LOO	0.327 (0.511)	0.880 (0.644)	0.172	AR 95% CI [0.694, 1.067]
IV: CTOT	−0.510 (0.978)	1.278 (1.737)	0.462	
IV: Joint (LOO+CTOT)	0.067 (0.529)	0.925 (0.646)	0.152	Hansen $J = 1.384$, $p = 0.501$
Panel B: Joint $\theta(D)$ at stability tertile centers				
Low ($D < -0.18$)	$\theta \approx 0$			
Mid ($D \approx 0.30$)	$\theta \approx 0.34$			
High ($D > 0.75$)	$\theta \approx 0.76$			

The institutional gradient is a structural feature of the data. Below the bottom tertile of political stability, the trade-network growth channel does not operate. Above the top tertile, countries absorb roughly 76% of partner growth shocks one-for-one. The mechanism is internal: own-country institutions gate the absorption channel, and the gradient is monotone, smooth, and economically large.

Figure 3 closes Section 5 by showing a comparison of each estimation across the full domain of Political Stability indicators.

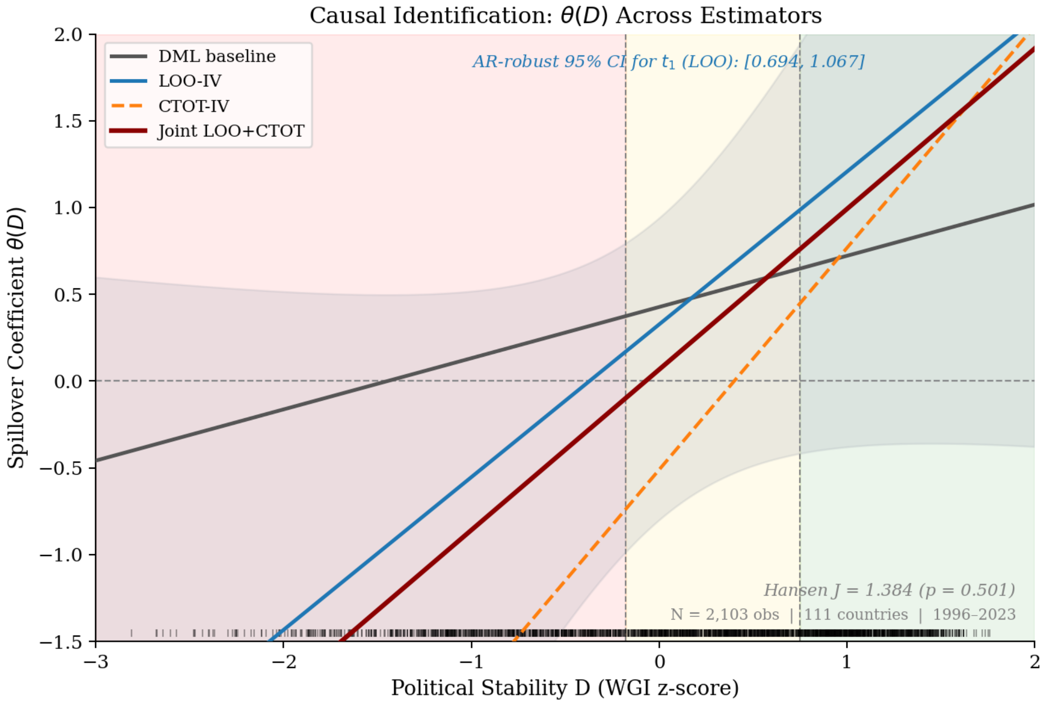

Figure 3: **Causal Identification: $\theta(D)$ Across Estimators.** Comparison of the implied institutional moderating function $\theta(D) = t_0 + t_1 \cdot D$ across the DML baseline, the LOO-IV, the CTOT-IV, and the joint LOO+CTOT estimators. The shaded band represents the 90% confidence interval for the joint estimator. All three IV specifications steepen the gradient relative to the DML baseline, with the joint estimator yielding $t_1 = 0.925$. The Hansen J test fails to reject the consistency of both instruments ($J = 1.384$, $p = 0.501$), and the Anderson-Rubin weak-IV-robust 95% confidence interval for t_1 from the LOO specification is $[0.694, 1.067]$, lying entirely above the DML baseline estimate of 0.295.

Discussion

The central finding of this paper is that the absorption rate of partner growth is not uniform across countries. It is a smooth, monotone function of own-country political stability, ranging from approximately zero in the bottom tertile of the stability distribution to above one in the top tertile. The economic interpretation of this gradient is straightforward and aligns with intuition: institutional stability gates the capacity to convert external opportunity into domestic growth. When partner growth rises, demand for i's exports rises along with it. Whether i's firms expand capacity to meet that demand depends on whether they can secure financing, enforce contracts, hire workers, and invest in fixed capital without political risk. Stable institutions make all of these decisions easier, and the frictions associated with unstable institutions make it either harder or downright impossible.

The High-tertile result of $\theta > 1$ deserves its own emphasis. A coefficient above one means that stable countries do not simply absorb partner growth one-for-one, they amplify it. The cleanest interpretation is a financial-accelerator mechanism. When external demand rises and the institutional environment supports leverage, firms borrow against expected revenue and expand capacity beyond the proportional response. Reliable contract enforcement, deep credit markets, and predictable trade policy are the conditions under which this amplification is possible. The Low-tertile result is the mirror image. In the absence of basic political stability, none of these channels operate. Capital flees rather than expands, contracts cannot be enforced, and firms have no incentive to invest in capacity to meet a demand surge they cannot reliably serve. The trade-network growth channel is closed for unstable countries.

These results have implications for how the existing trade-spillover literature should be read. The standard specification, exemplified by Arora and Vamvakidis (2005) and the broader IMF spillover literature, regresses domestic growth on a trade-weighted partner growth term and reports a single absorption coefficient. My results show that this coefficient is a sample-weighted average of a function $\theta(D)$ that varies by an order of magnitude across the institutional distribution. Two consequences follow. First, cross-country comparisons of "trade exposure" or "openness benefits" are implicitly conditional on the sample's institutional composition. Studies using mostly OECD countries will report large absorption coefficients. In contrast, studies including more emerging markets will report smaller ones. The apparent disagreement in the literature on the magnitude of trade spillovers may partly reflect this. Second, policy advice derived from these models is incomplete. The advice should be conditioned on institutional capacity, because for unstable countries the institutional bottleneck binds before the trade margin does.

This paper is closest in spirit to Coe et al. (2009), who first documented that institutional quality moderates international R&D spillovers. My contribution extends their work along three dimensions. I study growth spillovers rather than R&D spillovers, which speaks to a broader and more macroeconomically central channel. I allow flexible nonlinear $\theta(D)$ rather than imposing a linear interaction, which the polynomial robustness checks confirm is the correct functional form here. And I introduce an instrumental variable apparatus that separates the causal moderation effect from confounding through observables, which their original framework did not address.

The policy implications follow directly from the gradient itself. Trade-led development strategies require an institutional precondition. The Low-tertile finding implies that opening to trade in the absence of basic political stability does not generate the growth spillovers the strategy assumes. This challenges the implicit sequencing in some development prescriptions, where trade liberalization is treated as a stand-alone reform. My results suggest that without institutional stability, the trade channel does not operate, and resources spent on liberalization without institutional reform may have little-to-no effect. A second implication is that the return to institutional reform is multiplicative rather than additive. Improving political stability raises growth directly, as the existing institutions-and-growth literature has long established, but it also raises the rate at which a country captures growth from its trade partners. The total return is larger than the direct effect alone.

A useful illustration of the mechanism is Ukraine's institutional collapse following 2014. Applying the panel-estimated gradient to Ukraine's actual stability trajectory, the model implies that had Ukraine maintained the institutional level of a Central European peer such as Poland over 2014 to 2023, it would have absorbed roughly 15 additional percentage points of cumulative growth from partner networks over the decade. This is not a causal claim. It is an accounting exercise that applies a cross-country gradient to a single time series. The point is qualitative. The institutional shock did not just lower Ukraine's growth directly through the well-known channels of war and uncertainty, it also severed the country's connection to the global growth network entirely. Even when partner growth was strong, Ukraine's implied θ was near zero or negative.

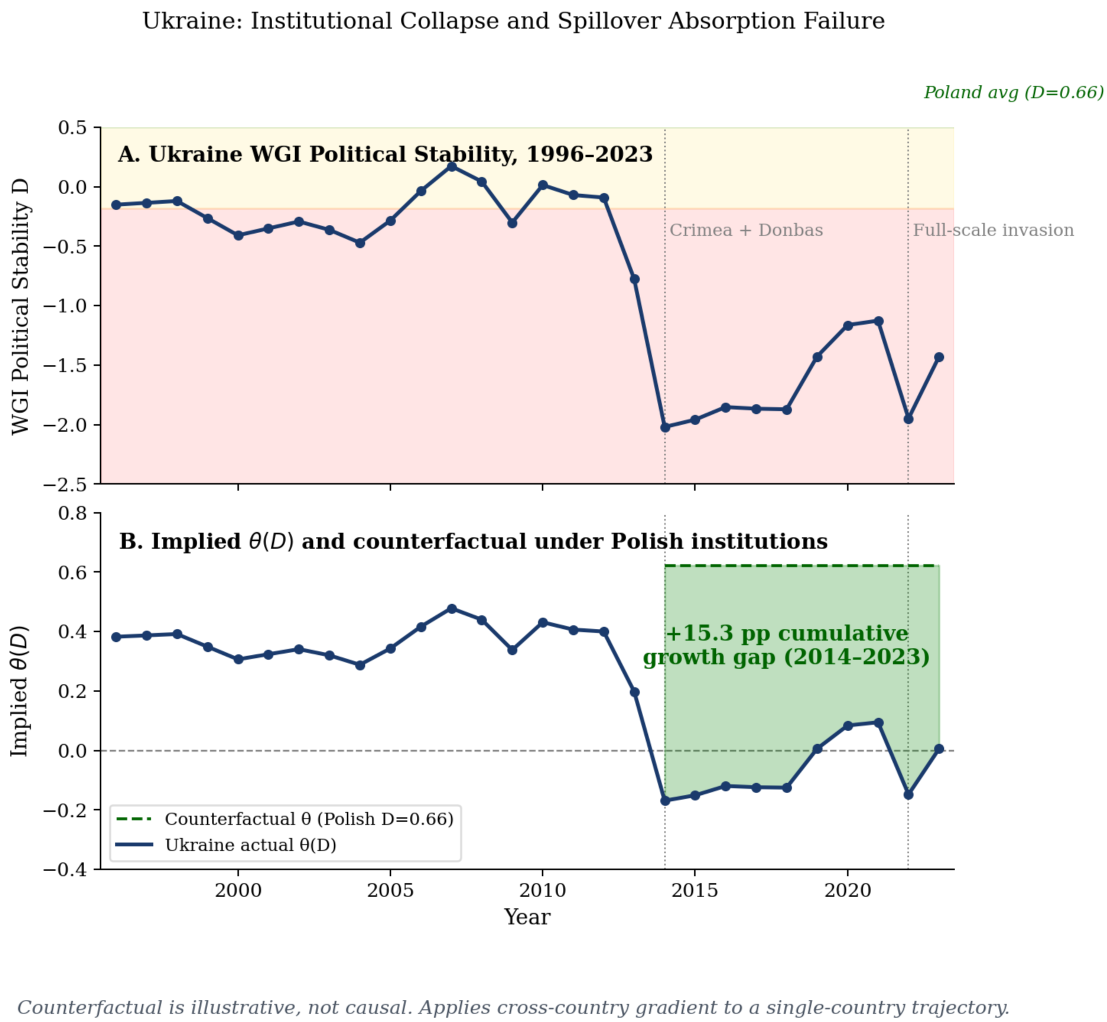

Figure 4: **Ukraine: Institutional Collapse and Spillover Absorption Failure.** Panel A shows Ukraine's WGI Political Stability trajectory from 1996 to 2023, with horizontal shading marking the Stable, Average, and Unstable tertile regimes. The dashed green line marks Poland's average D from 2014 to 2023. Panel B shows the implied $\theta_t = 0.427 + 0.295 \cdot D_{UKR,t}$ alongside the counterfactual $\theta_{CF} = 0.622$ that would have applied under Polish institutional levels. The shaded green region represents the cumulative absorption capacity lost over the decade. Applying the panel-estimated gradient to Ukraine as an accounting exercise, the model implies a counterfactual cumulative growth gain of approximately 15 percentage points had Ukraine maintained Polish-level institutional stability. This is illustrative rather than causal.

Two limitations deserve mention. First, the sample begins in 1996 because of WGI data availability. Robustness checks using V-Dem indicators that extend further back confirm the gradient survives, but the primary specification is restricted to the post-1996 period. Second, my results identify the absorption gradient through trade-weighted spillovers but do not isolate which specific institutional features drive the moderation. The robustness across alternative governance dimensions suggests the result is a feature of institutional capacity broadly rather than any single dimension, but a more granular decomposition of which institutional channels matter most is left for future work.

Robustness Checks

7.1 Functional Form of $\theta(D)$

A natural concern with the linear specification $\theta(D) = t_0 + t_1 D$ is that it imposes structure on a relationship that may have curvature. The tertile pattern, with θ rising sharply from Low to Mid and more gradually from Mid to High, hints visually at diminishing returns to institutional quality. I test this directly by augmenting the second stage with D^2 and D^3 terms.

The quadratic coefficient t_2 enters at +0.016 (SE = 0.100, $p = 0.871$), and the formal test does not reject the linear specification. The implied vertex of the quadratic, at $D^* = -t_1/(2t_2) = -9.118$, falls far outside the observed data range of $[-3, +2]$, confirming that the function is monotone over all observed institutional levels. Adding a cubic term yields the same conclusion. The cubic coefficient $t_3 = -0.068$ ($p = 0.333$) is insignificant, and t_1 remains positive and significant at the 5% level under cubic augmentation.

This implies that the linear specification is the appropriate model. The institutional moderating effect is smooth and progressive, not a regime switch or threshold. Each one standard deviation of political stability adds roughly 0.30 percentage points to the spillover elasticity, and this gradient is continuous through the high end of the distribution. The tertile pattern that visually suggested diminishing returns is an artifact of the arbitrary cutpoints, not a feature of the underlying function.

7.2 Alternative Moderators

The choice of WGI Political Stability as the primary moderator D raises the question of whether the gradient is specific to that single dimension or whether it reflects a broader feature of governance quality. To address this, I re-run the full DML pipeline with each of the remaining five WGI dimensions as alternative moderators, holding the rest of the specification fixed.

The slope t_1 is positive and significant at the 10% level in all six cases, and significant at the 1% level in five. The estimated t_1 ranges from 0.221 (WGI Rule of Law, $p = 0.054$) to 0.339 (WGI Voice and Accountability). The High-tertile θ is significant at the 1% level for every measure, ranging from 0.76 to 0.94 across the seven specifications. The Low-tertile θ is statistically indistinguishable from zero in five of the six specifications.

The institutional moderating gradient is a structural feature of governance quality broadly, not an artifact of the political stability dimension. WGI Regulatory Quality and Voice and Accountability yield the largest t_1 estimates, suggesting that business regulation and democratic accountability may matter slightly more for the trade-network channel than pure conflict and stability, but the differences are within the range of standard error overlap. The result that survives across all available measures is the structural one: countries with higher institutional quality absorb meaningfully more growth spillover from trade partners, and countries below a baseline level of governance are disconnected from the network entirely.

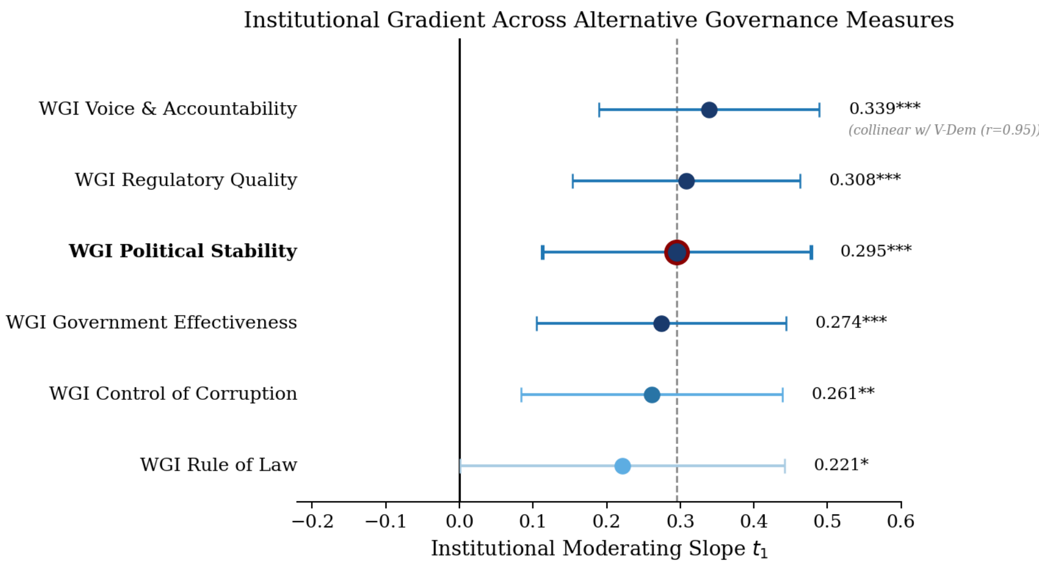

Figure 5: **Institutional Gradient Across Alternative Governance Measures.** Estimated slope t_1 from the linear DML specification $\theta(D) = t_0 + t_1 \cdot D$ across five alternative measures of institutional quality, with 90% confidence intervals. Each estimate uses the identical DML pipeline (5-fold Ridge cross-fitting, two-way fixed effects, country-clustered standard errors) and substitutes only the moderator variable. The baseline specification (WGI Political Stability) is highlighted in dark red. All six indicators yield positive t_1, with six significant at the 10% level or better.

7.3 First-Stage Method Robustness

A separate concern is whether the result depends on the choice of nuisance learner in the DML first stage. To address this, I re-estimate the baseline DML with four alternative learners: classical OLS, Ridge, Lasso, and Random Forest. The point estimate of t_1 ranges from 0.284 to 0.295 across all four methods, with significance preserved at the 1% level throughout. The result is not driven by any specific choice of first-stage estimator.

I also compare the cross-fitted Ridge residuals against classical Frisch-Waugh-Lovell OLS residuals directly. The two residual series correlate at $r = 0.999$ for both the treatment and outcome, and the resulting point estimates of t_1 differ by 1.3%. In this setting, DML and FWL-OLS produce numerically equivalent answers. The contribution of the DML framework is not a different point estimate but valid inference under potential nonlinearity in $f(X)$, and robustness to first-stage misspecification more generally. The substantive finding survives whether or not the reader accepts the ML machinery.

7.4 Lagged Growth Proxy Check

A final concern is that WGI Political Stability may proxy for recent growth performance rather than a genuine institutional channel. If stable countries are simply on faster growth trajectories, the estimated gradient t_1 would be misleading. I address this in two ways.

First, I regress D on lagged growth after two-way fixed-effect removal. The resulting R^2 is 0.0015, with raw correlations near zero. Within-country variation in political stability is near-orthogonal to lagged growth, ruling out the proxy story directly.

Second, I run an augmented DML specification that includes lag-1 GDP per capita growth as a direct first-stage control. This absorbs any residual growth momentum before estimating $\theta(D)$. The slope attenuates from 0.295 to 0.210 under Ridge and remains significant at the 5% level ($p = 0.041$). Across alternative first-stage learners, t_1 ranges from 0.146 to 0.210 in the augmented specification, with the lower bound coming from Random Forest. The conservative range $[0.146, 0.210]$ is uniformly above zero, and the threshold structure across tertiles is unaffected by this check since it does not rely on the linear slope.

The baseline $t_1 = 0.295$ is therefore a slight overestimate. The true gradient is plausibly closer to 0.21, but the qualitative conclusion is unchanged: institutional stability raises the absorption rate of partner growth shocks, and this is not an artifact of D proxying for growth momentum.

7.5 Placebo and Exclusion Diagnostics

The DML and IV results above rest on two implicit claims: that the institutional gradient operates through the trade-transmission channel, and that each instrument satisfies its exclusion restriction. This section reports placebo and falsification tests against both.

I first test whether the institutional gradient is specific to the trade-transmission horizon by replacing the spillover term $\tilde{g}_{i,t}$ with its k-period lag $\tilde{g}_{i,t-k}$ for $k \in \{0, 1, 2, 3, 4\}$, holding the outcome $g_{i,t}$ and moderator $D_{i,t}$ at time t. Figure 6 reports the resulting t_1 estimates. The slope decays monotonically from 0.291 at the baseline ($k = 0$, $p = 0.008$) to -0.020 at $k = 4$ ($p = 0.890$), with the signal collapsing by $k = 3$ ($p = 0.804$). The institutional gradient operates at the trade-transmission horizon, not through a persistent lag-invariant relationship between political stability and domestic growth.

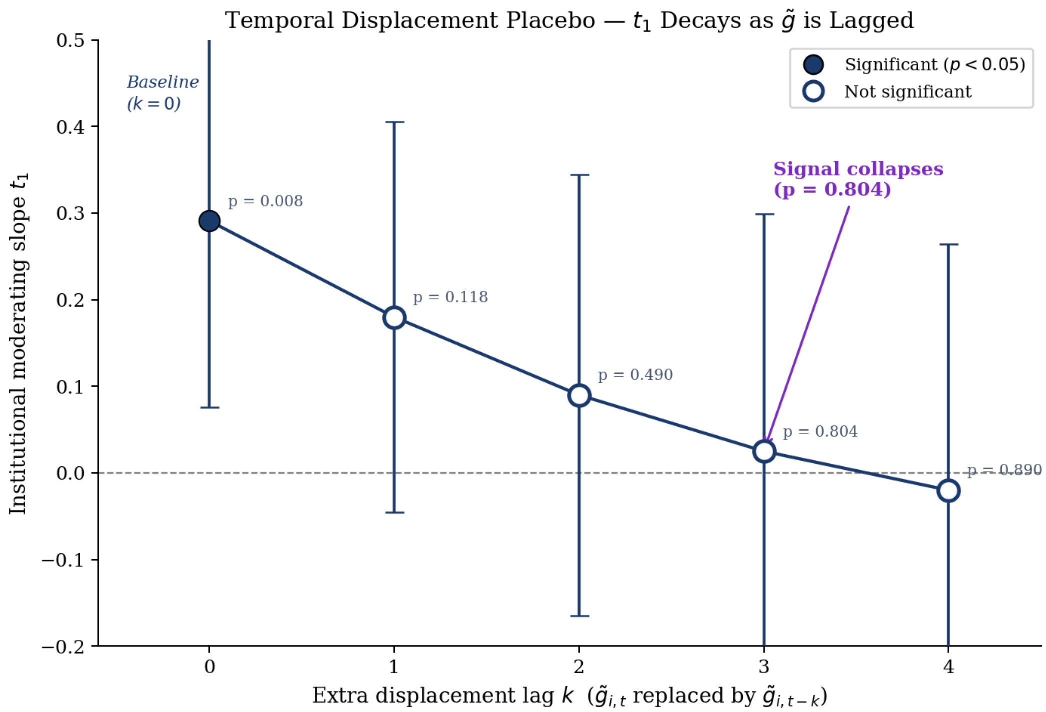

Figure 6: **Temporal Displacement Placebo: t_1 Decays as $\tilde{g}$ is Lagged.** The institutional moderating slope t_1 from the baseline DML specification, re-estimated with the trade-weighted spillover term $\tilde{g}_{i,t}$ replaced by its k-period lag $\tilde{g}_{i,t-k}$ for $k \in \{0, 1, 2, 3, 4\}$. The outcome $g_{i,t}$ and moderator $D_{i,t}$ are held at time t. Filled marker indicates significance at the 5% level; hollow markers are not significant. Error bars are 95% cluster-robust confidence intervals. The slope decays monotonically from $t_1 = 0.291$ ($p = 0.008$) at the baseline ($k = 0$) to $t_1 = -0.020$ ($p = 0.890$) at $k = 4$, with the signal collapsing by $k = 3$ ($p = 0.804$). The moderating effect operates at the trade-transmission horizon and not through a persistent lag-invariant relationship between institutional quality and domestic growth.

The Hansen J test in Section 5 provides indirect evidence for the exclusion restrictions by confirming that LOO and CTOT yield consistent estimates of $\theta(D)$. To support exclusion for each instrument individually, I run two additional placebo tests to address this concern.

The first is a lag-structure placebo on the CTOT instrument. After two-way fixed-effect removal, the correlation between Z^{CTOT}_{lag1} and own growth is -0.158. The sign is negative rather than positive, which is consistent with Dutch disease and mean-reversion after commodity booms. This is a substantive economic mechanism rather than a violation of exclusion. The exclusion restriction requires that Z^{CTOT} affect g only through $\tilde{g}$, not that the instrument be uncorrelated with g overall, and the Dutch disease channel operates orthogonally to the spillover channel.

The second is a direct-channel falsification on the variance decomposition of CTOT. Of the total CTOT variance, 13.9% is between-year and is fully absorbed by year fixed effects, leaving 49.1% as country-idiosyncratic variation. The IV identifies from the latter, which is the cleanest portion of the variation and is by construction orthogonal to the global commodity super-cycle.

For the LOO instrument, I regress Z_{it}^{LOO} on lagged own growth $g_{i,t-1}$ after fixed-effect removal and find a near-zero coefficient. This confirms that the leave-out construction blocks own-growth backdoor channels into the instrument.

I also briefly note the failure of an alternative shift-share design. A standard Bartik instrument constructed from base-period trade weights and partner growth shocks fails the Hansen J test against LOO and CTOT ($J = 16.19$, $p < 0.001$). The failure is structural: the leave-out correction required for exclusion validity in network settings, established by Borusyak et al. (2025), is absent by construction in the standard Bartik design. This is a methodological note rather than a robustness failure of the main result. The Bartik analysis is reported for transparency.

7.6 Common Correlated Effects: A Methodological Note

A natural robustness check for multi-factor common shocks is the Common Correlated Effects estimator of Pesaran (2006), which augments the regression with cross-section averages $\bar{z}_t$ as proxies for unobserved global factors. In standard panel settings, CCE addresses the residual common-shock concern that two-way fixed effects do not fully resolve. In my setting, however, CCE is structurally inappropriate.

The reason is that the treatment $\tilde{g}_{it} = \sum_j W_{ij,t-1} \cdot g_{j,t-1}$ is itself a cross-sectional aggregation of partner growth. Under Assumption 4 of Pesaran and Tosetti (2011), CCE requires factor loading heterogeneity between individual observations and the cross-section average. This fails when bilateral trade weights are diffuse because the trade-weighted partner growth term and the equal-weighted cross-section average converge to the same linear combination of common factors as the cross section grows. Empirically, the correlation between $\bar{z}(\tilde{g}_t)$ and the FE-demeaned treatment is 0.858 in my sample. Including the cross-section average as a control absorbs 74% of treatment variance, with the first-stage R^2 for the spillover term rising from 0.015 in the baseline DML to 0.738 under CCE. This collapses identification of $\theta(D)$ entirely. Excluding it leaves common shocks uncontrolled.

The LOO instrument resolves this tension by construction. It uses bilateral variation that is orthogonal to any equal-weighted cross-section average, which makes it the appropriate common-factor control for network-structured treatments. The CCE diagnostic is therefore not a robustness failure of the main result but rather a methodological finding about the limits of factor-proxy methods when the treatment variable is itself defined as a weighted cross-sectional aggregation.

Conclusion

This paper documents that the absorption rate of partner growth across bilateral trade networks is not uniform but is moderated by own-country political stability. The DML baseline establishes a positive and significant gradient. Causal identification through the

Leave-i-Out and Commodity Terms of Trade instruments confirms that this gradient is in fact substantially larger under exogenous variation, with the joint estimator placing t_1 near 0.93 and the Hansen J test failing to reject the consistency of both instruments. The implied function $\theta(D)$ ranges from approximately zero in the bottom tertile of the stability distribution to above one in the top tertile, with a smooth and monotone gradient connecting them.

The result reframes how the existing trade-spillover literature should be read. Specifications that impose a uniform absorption coefficient on the trade-weighted partner growth term are estimating a sample-weighted average of a function that varies by an order of magnitude across the institutional distribution. Some of the variation in the literature on the magnitude of trade spillovers may reflect differences in sample composition rather than genuine differences in the underlying mechanism, and policy advice derived from these models should be conditioned on the institutional capacity of the country in question.

Several extensions are natural. A more granular decomposition of which institutional channels drive the absorption gradient is the most direct continuation, and the alternative-D robustness in Section 7 suggests this exercise is feasible across the WGI and V-Dem dimensions. A second extension is to study whether the gradient varies with the type of partner growth shock, separating commodity-driven booms from technology-driven productivity gains. A third is to apply the same DML-IV apparatus to other network-structured treatments where standard tools such as CCE and Bartik fail in similar ways. The methodological contribution about network treatments may extend further than the substantive finding about institutions and growth.

References

Abramovitz, Moses, "Catching Up, Forging Ahead, and Falling Behind," *The Journal of Economic History*, 1986, *46* (2), 385–406.

Acemoglu, Daron, Simon Johnson, and James A. Robinson, "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, December 2001, *91* (5), 1369–1401.

_ , _ , and _ , "Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution," *The Quarterly Journal of Economics*, November 2002, *117* (4), 1231–1294.

_ , _ , _ , and Yunyong Thaicharoen, "Institutional causes, macroeconomic symptoms: volatility, crises and growth," *Journal of Monetary Economics*, January 2003, *50* (1), 49–123.

_ , Ufuk Akcigit, and William Kerr, "Networks and the Macroeconomy: An Empirical Exploration," *NBER Macroeconomics Annual*, January 2016, *30* (1), 273–335.

Adler, Gustavo and Nicolas E Magud, "Commodity price shocks and GDP growth in developing countries: The role of the exchange rate regime," *IMF Economic Review*, 2015, *63* (4), 863–891.

Aisen, Ari and Francisco José Veiga, "How Does Political Instability Affect Economic Growth?," *IMF Working Papers*, 2011, *11* (12), 1.

Alesina, Alberto, Sule Zler, Nouriel Roubini, and Phillip Swagel, "Political instability and economic growth," *Journal of Economic Growth*, June 1996, *1* (2), 189–211.

Alfaro, Laura, Areendam Chanda, Sebnem Kalemli-Ozcan, and Selin Sayek, "FDI and economic growth: the role of local financial markets," *Journal of International Economics*, 2004, *64* (1), 89–112.

Arora, Vivek and Athanasios Vamvakidis, "How Much Do Trading Partners Matter for Economic Growth?," *IMF Staff Papers*, 2005, *52* (1), 24–40.

Barro, Robert J., "Economic Growth in a Cross Section of Countries," *The Quarterly Journal of Economics*, May 1991, *106* (2), 407.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen, "High-dimensional methods and inference on structural and treatment effects," *Journal of Economic Perspectives*, 2014, *28* (2), 29–50.

Borensztein, E., J. De Gregorio, and J.-W. Lee, "How does foreign direct investment affect economic growth?1," *Journal of International Economics*, 1998, *45* (1), 115–135.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "A Practical Guide to Shift-Share Instruments," *Journal of Economic Perspectives*, February 2025, *39* (1), 181–204.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin, "Identification of peer effects through social networks," *Journal of Econometrics*, 2009, *150* (1), 41–55.

Cameron, A. Colin and Douglas L. Miller, "A Practitioner's Guide to Cluster-Robust Inference," *Journal of Human Resources*, 2015, *50* (2), 317–372.

Caselli, Francesco, Miklós Koren, Milan Lisicky, and Silvana Tenreyro, "Diversification through Trade," *The Quarterly Journal of Economics*, 2020, *135* (1), 449–502.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins, "Double/debiased machine learning for treatment and structural parameters," *The Econometrics Journal*, February 2018, *21* (1), C1–C68.

Clarke, Paul S and Annalivia Polselli, "Double machine learning for static panel models with fixed effects," *The Econometrics Journal*, January 2026, *29* (1), 69–86.

Coe, David T. and Elhanan Helpman, "International R&D spillovers," *European Economic Review*, May 1995, *39* (5), 859–887.

_ , _ , and Alexander W. Hoffmaister, "International R&D spillovers and institutions," *European Economic Review*, October 2009, *53* (7), 723–741.

Cohen, Wesley M. and Daniel A. Levinthal, "Absorptive Capacity: A New Perspective on Learning and Innovation," *Administrative Science Quarterly*, March 1990, *35* (1), 128.

Coppedge, Michael, John Gerring, Carl Henrik Knutsen, Staffan I. Lindberg, Jan Teorell et al., "V-Dem Dataset v14," 2024.

Deaton, Angus and R.-I. Miller, "International Commodity Prices, Macroeconomic Performance, and Politics in Sub-Saharan Africa," Working paper, International Economics Section, Departement of Economics Princeton University, 1995.

Drechsel, Thomas and Silvana Tenreyro, "Commodity booms and busts in emerging economies," *Journal of International Economics*, May 2018, *112*, 200–218.

Durham, J. B. J. Benson, "Absorptive capacity and the effects of foreign direct investment and equity foreign portfolio investment on economic growth," *European Economic Review*, 2004, *48* (2), 285–306.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, October 2015, *105* (10), 3150–3182.

Fernández, Andrés, Stephanie Schmitt-Grohé, and Martín Uribe, "World shocks, induced business cycles, and emerging economies," *Journal of International Economics*, 2017, *108*, 1–15.

Feyrer, James, "Trade and Income—Exploiting Time Series in Geography," *American Economic Journal: Applied Economics*, October 2019, *11* (4), 1–35.

Frankel, Jeffrey A. and David H. Romer, "Does Trade Cause Growth?," *American Economic Review*, June 1999, *89* (3), 379–399.

Giovanni, Julian Di and Andrei A Levchenko, "Putting the Parts Together: Trade, Vertical Linkages, and Business Cycle Comovement," *American Economic Journal: Macroeconomics*, April 2010, *2* (2), 95–124.

Glaeser, Edward, Rafael La Porta, Florencio Lopez de Silane, and Andrei Shleifer, "Do Institutions Cause Growth?," Technical Report w10568, National Bureau of Economic Research, Cambridge, MA June 2004.

Hall, R. E. and C. I. Jones, "Why do Some Countries Produce So Much More Output Per Worker than Others?," *The Quarterly Journal of Economics*, February 1999, *114* (1), 83–116.

International Monetary Fund, "Direction of Trade Statistics (DOTS)," https://data.imf.org/dots 2024.

Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi, *The worldwide governance indicators : methodology and analytical issues*, World Bank, September 2010.

Kebhaj, Suhaib, "A new monthly index of commodity terms of trade," IMF Working Paper WP/20/12, International Monetary Fund 2020.

Keller, Wolfgang, "International Technology Diffusion," *Journal of Economic Literature*, August 2004, *42* (3), 752–782.

Manski, Charles F., "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies*, July 1993, *60* (3), 531.

Marini, Marco, Robert Dippelsman, and Michael Stanger, "New Estimates for Direction of Trade Statistics," *IMF Working Papers*, 2018, *18*, 1.

Pemstein, Daniel, Kyle Marquardt, Eitan Tzelgov, Yi ting Wang, Juraj Medzihorsky, Joshua Krusell, Farhad Miri, and Johannes von Römer, "The V-Dem Measurement Model: Latent Variable Analysis for Cross-National and Cross-Temporal Expert-Coded Data," Technical Report Working Paper No. 21, 10th edition, Varieties of Democracy Institute, University of Gothenburg 2025.

Pesaran, M. Hashem, "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure," *Econometrica*, July 2006, *74* (4), 967–1012.

_ **and Elisa Tosetti**, "Large panels with common factors and spatial correlation," *Journal of Econometrics*, April 2011, *161* (2), 182–202.

Rodrik, Dani, "Where Did All the Growth Go? External Shocks, Social Conflict, and Growth Collapses," *Journal of Economic Growth*, December 1999, *4* (4), 385–412.

_ **, Arvind Subramanian, and Francesco Trebbi**, "Institutions Rule: The Primacy of Institutions over Geography and Integration in Economic Development," Technical Report w9305, National Bureau of Economic Research, Cambridge, MA 11 2002.

Semenova, Vira and Victor Chernozhukov, "Debiased Machine Learning of Conditional Average Treatment Effects and Other Causal Functions," 2017. Version Number: 6.

Wacziarg, Romain and Karen Horn Welch, "Trade Liberalization and Growth: New Evidence," *The World Bank Economic Review*, January 2008, *22* (2), 187–231.