FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Financial Conditions and Capital Investment Choices

Òscar Jordà
Federal Reserve Bank of San Francisco
University of California, Davis and CEPR

Fernanda Nechio
Federal Reserve Bank of San Francisco

Toan Phan
Federal Reserve Bank of Richmond

Felipe Schwartzman
Federal Reserve Bank of Richmond

February 2026

Working Paper 2026-05

https://doi.org/10.24148/wp2026-05

Suggested citation:

Jordà, Òscar, Fernanda Nechio, Toan Phan, and Felipe Schwartzman. 2026. "Financial Conditions and Capital Investment Choices." Federal Reserve Bank of San Francisco Working Paper 2026-05. https://doi.org/10.24148/wp2026-05

Financial conditions and capital investment choices*

Òscar Jordà[†] Fernanda Nechio[‡] Toan Phan[§]
Felipe Schwartzman[¶]

26th February 2026

Abstract

We show, both theoretically and empirically, that tight financial conditions shift investment toward cheaper but less energy-efficient capital. In a small open-economy model with vintage capital, higher financing costs reduce the present value of future energy savings, tilting firms' choices along a cost-efficiency frontier. Using 150 years of macroeconomic and energy data from 17 advanced economies, we find that tighter financial conditions reduce output, capital, and total energy consumption, but raise the amount of energy per unit of capital (energy intensity), a composition effect that persists for 6 to 8 years. Tight financial conditions lower energy use in the short run by depressing activity, but increase energy use in the medium run through worse energy efficiency.

Keywords: energy efficiency, capital vintages, monetary policy, interest rates, local projections, small open economy

*The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Richmond, the Federal Reserve Bank of San Francisco, or the Federal Reserve System. We thank Brooke Hansbrough and River Bell for excellent research assistance. All errors are our own.

[†]Federal Reserve Bank of San Francisco; and Department of Economics, University of California, Davis; and CEPR (oscar.jorda@sf.frb.org; ojorda@ucdavis.edu).

[‡]Federal Reserve Bank of San Francisco (fernanda.nechio@sf.frb.org).

[§]Federal Reserve Bank of Richmond (toan.phan@rich.frb.org).

[¶]Federal Reserve Bank of Richmond (felipe.schwartzman@rich.frb.org).

1. Introduction

Traditional analysis of monetary economies centers around the effects of financial conditions on inflation and the business cycle (among the many textbook treatments, see, *e.g.*, Galí, 2015). However, recent research suggests that financial conditions can also have longer-run effects on the economy and in particular, on investment and total factor productivity (TFP) growth (see, *e.g.*, Bernanke and Mihov, 1998; Brunnermeier et al., 2021; Jordà et al., 2024). As investment in new capital embodies choices about energy consumption, and since these choices depend on financial costs, it is natural to ask how financial conditions affect the composition of investment and, in particular, the intensity of energy use.

On one hand, a loose financial stance boosts overall economic activity and energy use. On the other hand, a loose stance lowers the discount rate applied to future energy savings, encouraging firms to invest in costlier but more energy-efficient capital. This tension is the central question that this paper investigates. A key challenge in answering this question is that the relevant adjustment margin operates through the slow turnover of the capital stock rather than at business-cycle frequencies. Energy-efficiency choices are embedded in long-lived capital vintages, so their aggregate consequences unfold over decades. This makes short samples ill-suited to detect the composition effects implied by the theory and motivates our use of a long historical panel spanning 150 years.

Our main finding is consistent with a *composition effect*: tighter financial conditions reduce total energy use by depressing investment and output, but simultaneously shift the composition of new investment toward less energy-efficient capital. Energy intensity of capital (*i.e.*, the energy use per unit of capital) *rises* even as total energy consumption falls. This composition effect is persistent, lasting 6 to 8 years, and reflects the slow turnover of the capital stock. We establish this result in both a tractable model and in a panel data analysis covering 17 advanced economies over 150 years.

The model is a small open economy with vintage capital in which a representative planner chooses the energy intensity of each new cohort of investment along a technological frontier that trades off up-front capital cost against ongoing energy requirements. The key model equation is simple: optimal energy intensity equates the marginal rate of transformation along this frontier to the present value of future energy costs, discounted at the world interest rate and adjusted for depreciation. Higher financing costs reduce this present value, making it optimal to select cheaper but less efficient capital. Because the installed capital stock is a mix of overlapping vintages, the effect on average energy intensity is hump-shaped—muted initially by the presence of older, less efficient vintages, then rising as these depreciate.

The vintage model implies slow adjustment. Empirically analyzing prolonged periods of loose/tight financial conditions and their medium- to long-run effects requires a long time series and a cross section of countries to extract the necessary statistical power. We combine the Jordà et al. (2017) historical macrofinancial database with historical data on energy use and data on capital from the Long-Term Productivity Database.[1] We enhance our sample by also using data on CO2 emissions from the Our World in Data (OWID) project, which correlate well with data on energy consumption, and is available for a much longer time sample and cross-section. Therefore, in our empirical analysis we either rely on energy consumption data for a shorter time sample, or proxy energy consumption with data on CO2 emissions to obtain a longer time series.

We measure financial conditions by using the exogenous variation in the stance of monetary conditions. Establishing whether the stance of financial conditions is loose or tight naturally requires a benchmark. Our approach is to borrow directly from the literature on monetary policy and obtain estimates of the natural rate of interest. For the latter, recent work by Del Negro et al. (2019) provides estimates of the natural rate based on 7 advanced economies, which Grimm et al. (2023) extend to the set of advanced economies used in our analysis. We also borrow from the latter a definition for the monetary stance which we detail more carefully below.

Of course, fluctuations in the monetary stance probably respond to conditions in the economy that could also relate to energy use. This requires finding a source of exogenous variation in the monetary stance to allow for a causal interpretation of the results. Our approach borrows from the *trilemma* instrumental variable approach introduced in Jordà et al. (2020). The logic connects directly to the model: in a small open economy that pegs its exchange rate and capital moves freely, domestic interest rates are driven by the base-country financial conditions rather than domestic considerations. The trilemma instrument exploits this by isolating the component of base-country interest rate fluctuations that cannot be predicted by base-country macroeconomic variables, scaled by the degree of exchange rate pegging and capital openness. These spillovers provide a source of exogenous variation in the domestic financial stance that is plausibly orthogonal to domestic energy-market conditions.

Methodologically, identifying composition effects driven by slow capital turnover requires an empirical approach that is robust to long lags and flexible about the timing and shape of adjustment. We therefore use local projections (Jordà, 2005), which are well suited to tracing medium- and long-run responses when capital adjustment is gradual and per-

[1] For details on the database and project see Bergeaud et al. (2016).

sistent. Local projections are also less sensitive than VARs to lag-length and specification choices for slow-moving variables (Olea et al., 2024). Using this framework, we estimate how output, energy use, and energy intensity respond to fluctuations in monetary conditions over time. More specifically, our empirical analysis implements local projections using instrumental variables (Jordà and Taylor, 2016; Plagborg-Møller and Wolf, 2021) on panel data.

Our findings line up well with intuition. More importantly, they are also in line with our theoretical predictions. Tighter monetary conditions are associated with lower output and energy consumption. However, at the same time, investment falls and energy intensity of capital increases. This is particularly the case when energy intensity is measured relative to "machine" capital rather than "buildings" capital, consistent with the model's intuition that the relevant margin of adjustment is in equipment–the type of capital where energy-efficiency choices are most salient. At the same time, we do not find evidence of permanent changes in energy efficiency of the capital stock of the economy, consistent with changing choices of technology within a given menu that is reversible as conditions normalize.

Related literature. Our paper contributes to three strands of literature. First, we contribute to the literature that studies long-run real effects of monetary policy, which includes Blanchard and Summers (1986); Bernanke and Mihov (1998); Garga and Singh (2021), among others. Most related, Jordà et al. (2024) show that contractionary monetary shocks reduce investment and TFP for extended periods. We extend this finding by showing that tight financial conditions affect not only the *quantity* of investment but also its *quality*—specifically, its energy efficiency. This quality margin is naturally interpreted through a vintage-capital lens in which energy use is tied to installed equipment in the short run and adjusts through the composition of new capital in the long run, as in the putty-clay tradition of Atkeson and Kehoe (1999).

Second, we relate to the literature on finance and the energy transition. Aghion et al. (2022) and Aghion et al. (2024) show that green R&D investment is particularly sensitive to financial conditions, while Aghion et al. (2019) document the sensitivity of R&D more broadly to monetary conditions. Känzig et al. (2025) find that green innovation is countercyclical—driven by a "green is in the future" channel whereby green patents derive value from expected future profits, making them less sensitive to short-term fluctuations. Also related, Känzig and Williamson (2025) study the macroeconomic drivers of energy-saving technical change and argue that energy-saving technology shocks are central to the decoupling of output and fossil energy use. Using Indian manufacturing microdata, McDonald et al. (2026) find a complementary channel through which energy intensity declines with growth: firms' energy expenditure shares fall steeply with scale, driven by physical scaling laws and fixed-cost efficiency investments. Our paper complements these papers by identifying a distinct channel that does not require R&D or a green/brown technology distinction: even in the absence of innovation, the *choice* of energy intensity from an existing technological frontier responds to interest rates through a present-value mechanism.

Third, our model builds on the energy-macroeconomic framework of Hassler et al. (2021), extending it with a vintage capital structure à-la Atkeson and Kehoe (1999) and an open-economy setting with exogenous interest rate fluctuations.[2] Within this literature, Casey (2024) develops a directed technical change model showing that the standard Cobb-Douglas specification overstates the energy reductions achievable through energy taxes and that exogenous efficiency improvements can paradoxically reduce incentives for further energy-saving R&D. Cakir Melek and Orak (2024) emphasize that time-varying substitutability between capital equipment and energy is central to understanding the declining energy share of income, highlighting the importance of equipment quality for energy efficiency. Amador (2022) studies technology adoption with fixed costs and externalities, showing that monetary shocks can induce reversion to older technologies. Our mechanism is complementary but more primitive: it operates on the continuous margin of energy-intensity choice within a cohort, without requiring fixed costs, externalities, discrete technology adoption, or R&D dynamics. The present-value channel we identify is robust to the specific form of the technological frontier—as we show, log-linear specifications with directed technical change in growth rates reduce to the same static tradeoff. Keuschnigg and Stalenis (2025) study energy-saving innovation and vintage capital in the green transition with endogenous scrapping and policy, highlighting how incentives shape the speed of diffusion across vintages; our contribution isolates a complementary forcediscount-rate variationthat shifts the composition of new vintages even absent innovation.

2. Model

To fix ideas, we first develop a small open economy model with capital vintages of different energy efficiencies. The model allows us to analyze how exogenous interest rate movements influence energy intensity (the amount of energy per unit of capital) through intertemporal investment choices. We build on the macro-energy framework established by Hassler et al. (2021), with two key departures. First, we allow each new *vintage* of cap-

[2] Also see Hassler et al. (2016) and Bilal and Stock (2024) for a survey of the recent environmental-macroeconomic literature.

ital to have a distinct *energy intensity*—energy used per unit of capital—that is chosen at the time of investment. For instance, if you think about an air-conditioning or a heating unit as a form of capital, a more energy-efficient model typically has a higher up-front cost but requires lower flows of energy – *i.e.*, lower energy intensity.[3] A higher interest rate environment effectively discounts those future cost savings more steeply, tilting firms towards less energy-efficient (but cheaper) capital. Second, to better map the model to the empirical context, we allow for capital flows that are subject to exogenous movements in the world interest rate.

The general intuition is as follows. When the world interest rate rises, future cost savings from energy-efficient capital are discounted more steeply, increasing the incentive to purchase less energy-efficient but cheaper capital. This effect raises *energy intensity* in the economy in the aftermath of the shock (even if total energy usage falls in absolute terms), with the persistence of the effect depending on the durability of capital.

2.1. Environment

Time is discrete, indexed by $t = 0, 1, 2, \ldots$. In each period, a representative planner chooses consumption C_t, external borrowing B_{t+1} (at an exogenous world interest rate $R_{t,t+1}$), energy use E_t (sold at an exogenous world energy price P_t), capital investment I_t, and *energy intensity* ϵ_t (from a technological frontier, to be described below).

Denote by $k_t(v)$ the amount of vintage-v capital remaining at time t. Investment I_t becomes *capital of vintage t* in the following period $(t+1)$ and has energy intensity $\epsilon(t) = \epsilon_t$. A uniform depreciation rate $\delta \in [0, 1]$ applies to all vintages. New investment I_{t-1} takes one period to become the active capital of vintage t.

$$k_t(v) = \begin{cases} (1-\delta)\, k_{t-1}(v), & v < t, \\ I_{t-1}, & v = t. \end{cases} \tag{1}$$

[3] Empirical evidence from Newell et al. (1999) suggests that firms respond to energy prices and regulations by shifting the available menu of durable capital good vintages toward more efficient designs, highlighting the same trade-off between initial investment and future energy costs present in our model.

Capital services and energy use. Each vintage v generates *capital services* $\tilde{k}_t(v)$ by combining capital $k_t(v)$ with energy input $e_t(v)$. For tractability, we assume a Leontief structure:[4]

$$\tilde{k}_t(v) \;=\; \min\{\, k_t(v),\; e_t(v)/\epsilon(v)\}.$$

Summing over vintages,

$$\tilde{K}_t \;=\; \sum_{v=0}^{t} \tilde{k}_t(v).$$

Cost minimization implies $e_t(v) = \epsilon(v)\, k_t(v)$: each unit of vintage-$v$ capital requires $\epsilon(v)$ units of energy to operate per period. Therefore, $\tilde{k}_t(v) = k_t(v)$ in equilibrium. Hence, total capital services are

$$\tilde{K}_t = \sum_{v=0}^{t} k_t(v).$$

Aggregate output is produced by combining capital services with inelastically supplied labor. Normalizing employment to 1, this induces the aggregate production function

$$Y_t = \left(\tilde{K}_t\right)^{\alpha} = \left(\sum_{v=0}^{t} k_t(v)\right)^{\alpha},$$

with $0 < \alpha < 1$. Total energy usage is

$$E_t \;=\; \sum_{v=0}^{t} e_t(v) \;=\; \sum_{v=0}^{t} \epsilon(v)\, k_t(v).$$

Using Equation 1, it follows that

$$E_t = \sum_{v=0}^{t} \epsilon(v)\, \left(1-\delta\right)^{\,t-v}\, I_{v-1}.$$

Hence, at the aggregate level, capital and energy follow simple dynamics:

$$K_{t+1} \;= (1-\delta)\, K_t \;+\; I_t, \tag{2a}$$

$$E_{t+1} \;= (1-\delta)\, E_t \;+\; \epsilon_t\, I_t. \tag{2b}$$

New investment I_t adds to the capital stock, while depreciation at rate δ reduces the

[4]This assumption reflects the empirical observation that, in the short run, firms have limited ability to substitute between capital and energy, making it reasonable to impose a low elasticity of substitution (León-Ledesma and Satchi, 2019; Hassler et al., 2021). In the long run, firms can adjust their technology mix, but over the short-to-medium horizons, substitution is constrained, as installed capital cannot quickly be retrofitted to change its energy requirements.

amount of usable capital. The key asymmetry is in Equation 2b: aggregate energy usage depends not only on the *quantity* of investment but also on its *energy intensity* ϵ_t, which is a choice variable. As a result, E_t is a weighted average across vintage-specific intensities, and shocks that shift ϵ_t permanently alter the composition of the installed capital stock.

The aggregate resource constraint is:

$$C_t + R_{t-1,t}\, B_t + q_t\, I_t + P_t\, E_t = B_{t+1} + K_t^{\alpha}. \tag{3}$$

Here, $R_{t-1,t}\, B_t$ is the gross repayment of external debt contracted in period $t-1$, B_{t+1} is new external borrowing at the current world interest rate $R_{t,t+1}$, and q_t is the up-front purchase cost of new capital, which depends on its energy intensity ϵ_t. The term $P_t E_t$ captures the flow cost of energy used in production.

Capital-energy technology frontier. The planner chooses sequences $\{C_t, I_t, E_t, B_{t+1}, \epsilon_t\}$ to maximize lifetime utility of a representative agent:

$$\max \sum_{t=0}^{\infty} \beta^t\, u(C_t),$$

subject to the resource constraint Equation 3, the laws of motion for capital and energy Equation 2a and Equation 2b, *and a technological frontier that constrains the relationship between capital cost and energy efficiency*:

$$G(\epsilon_t, q_t) \equiv \eta \log q_t + (1-\eta) \log \epsilon_t = 0, \quad 0 < \eta < 1. \tag{4}$$

This formulation captures the idea that more energy-efficient capital (lower ϵ_t) requires a higher up-front investment (higher q_t). This formulation of the frontier (which follows Jones 2005; León-Ledesma and Satchi 2019; Hassler et al. 2021) ensures that the planner (or agents in the decentralized economy) optimally weighs the immediate cost of investment against the long-run savings from lower energy expenditures. It is also consistent with empirical evidence from Newell et al. (1999), which shows that firms adjust the efficiency of available capital in response to changes in energy prices and regulatory policies, reinforcing the view that energy efficiency is an endogenous economic choice rather than an exogenous trend. The key parameter η governs the trade-off between capital cost q_t and energy intensity ϵ_t in our model:[5]

$$\frac{d \ln q_t}{d \ln \epsilon_t} = -\frac{1-\eta}{\eta}.$$

Mapping to the empirical setting. The model is written as a planner's problem for expositional clarity, but the solution coincides with the competitive equilibrium of a decentralized economy with price-taking firms and households facing the same world interest rate. The exogenous world interest rate $R_{t,t+1}$ in the model maps to exogenous part of the the empirical monetary stance identified by the trilemma instrument: in both model and data we focus on fluctuations in borrowing costs are exogenous to the individual economy. In the data, the trilemma instrumental variable isolates precisely this type of variation—interest rate movements driven by base-country monetary policy spillovers rather than domestic conditions.

The model's key endogenous variables have direct empirical counterparts. Output Y_t maps to real GDP, capital K_t to the real capital stock, and energy E_t to primary energy consumption (or, given the close empirical link between the two, CO2 emissions). The model's central prediction concerns the ratio E_t/K_t—average energy intensity of the installed capital stock—which corresponds to the E/K and CO2/K ratios in the empirical analysis. The energy intensity of new investment, ϵ_t, is not directly observed, but its aggregate effects are revealed through the dynamics of E_t/K_t.

2.2. Optimality

Let λ_t be the multiplier on Equation 3, χ_t and ξ_t be the multipliers enforcing Equation 2a and Equation 2b, respectively, and γ_t be the multiplier on $G(\epsilon_t, q_t) \geq 0$. The Lagrangian is

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \left\{ \begin{array}{cc} & u(C_t) \\ + & \lambda_t \big(K_t^\alpha + B_{t+1} - C_t - R_{t-1,t} B_t - q_t I_t - P_t E_t \big) \\ + & \chi_t \big((1-\delta) K_t + I_t - K_{t+1} \big) \\ - & \xi_t \big((1-\delta) E_t + \epsilon_t I_t - E_{t+1} \big) \\ + & \gamma_t G(\epsilon_t, q_t) \end{array} \right\}.$$

[5] This elasticity parallels Hassler et al. (2021), who estimate the trade-off between capital-augmenting technology growth g_A and energy-saving technology growth g_{A_e} as $\frac{d \ln g_A}{d \ln g_{A_e}} = -\frac{1}{\phi}$. They estimate from historical US data that $1/\phi \approx 11.5$, which implies a mild trade-off between capital cost and energy efficiency ($\eta \approx 1/(1+\phi) \approx 0.92$ in our model).

Taking first-order conditions:

$$\partial C_t: \qquad u'(C_t) \;=\; \lambda_t, \tag{5}$$

$$\partial B_{t+1}: \qquad \lambda_t \;=\; \beta\, R_{t,t+1}\, \lambda_{t+1}, \tag{6}$$

$$\partial K_{t+1}: \qquad \chi_t \;=\; \alpha\, \beta\, K_{t+1}^{\alpha-1}\, \lambda_{t+1} \;+\; (1-\delta)\, \beta\, \chi_{t+1}, \tag{7}$$

$$\partial E_{t+1}: \qquad \xi_t \;=\; \beta\, \lambda_{t+1}\, P_{t+1} \;+\; (1-\delta)\, \beta\, \xi_{t+1}, \tag{8}$$

$$\partial I_t: \qquad \chi_t \;=\; q_t\, \lambda_t \;+\; \epsilon_t\, \xi_t, \tag{9}$$

$$\partial \epsilon_t: \qquad G_\epsilon(\epsilon_t, q_t)\, \gamma_t \;=\; I_t\, \xi_t, \tag{10}$$

$$\partial q_t: \qquad G_q(\epsilon_t, q_t)\, \gamma_t \;=\; I_t\, \lambda_t. \tag{11}$$

Equation 5 and Equation 6 imply that the multiplier λ_t follows the usual Euler condition $\lambda_t = \beta\, R_{t,t+1}\, \lambda_{t+1}$. Equation 7 says that the shadow value of one extra unit of capital χ_t depends on the marginal product of capital (discounted by β), plus the value of the non-depreciated capital in the next period $(1-\delta)\, \chi_{t+1}$. Equation 8 has a symmetric form for the multiplier on energy ξ_t: it equals the discounted cost of one more unit of energy next period ($\beta\lambda_{t+1}P_{t+1}$) plus the continuation value of the energy obligation that persists on nondepreciated capital. Meanwhile, Equation 9 says the shadow value of capital χ_t equals the *purchase cost* $q_t\lambda_t$ plus the *energy cost* $\epsilon_t\xi_t$ of new investment. This decomposition is central: investing in a new vintage entails both an immediate capital expenditure and a commitment to future energy costs, with the latter depending on the chosen intensity ϵ_t.

Choice of energy efficiency ϵ_t. Equation 10 and Equation 11 pin down how ϵ_t and q_t move along the technology frontier $G(\epsilon_t, q_t) = 0$. Together they yield a condition linking ϵ_t and q_t to ξ_t/λ_t:

$$\frac{G_\epsilon(\epsilon_t, q_t)}{G_q(\epsilon_t, q_t)} = \frac{\xi_t}{\lambda_t} = \sum_{s=1}^{\infty} \frac{(1-\delta)^{s-1}}{R_{t,t+s}}\, P_{t+s}, \tag{12}$$

where the second equality emerges from iterating Equation 8 forward and using Equation 6 to substitute out future values of λ_{t+s}. Here $R_{t,t+s} \equiv \prod_{j=0}^{s-1} R_{t+j,t+j+1}$ denotes the compound gross interest rate from t to $t+s$.

The condition has a transparent interpretation. The left-hand side is the marginal rate of transformation along the technology frontier: how much additional capital cost q_t is required to achieve a unit reduction in energy intensity ϵ_t. The right-hand side is the present value of the stream of energy costs that one unit of capital imposes over its lifetime, discounted by both the interest rate and depreciation. The planner invests in energy efficiency up to the point where the marginal cost of further improvement (higher q_t) equals the discounted value of the energy savings it delivers.

For our functional form $G(\epsilon, q) = \eta \log q + (1-\eta)\log \epsilon = 0$, we have $G_\epsilon(\epsilon, q) = (1-\eta)/\epsilon$ and $G_q(\epsilon, q) = \eta/q$, so Equation 12 becomes:

$$\underbrace{\frac{1-\eta}{\eta}\frac{q_t}{\epsilon_t}}_{\text{"up-front" MRS}} = \underbrace{\sum_{s=1}^{\infty} \frac{(1-\delta)^{s-1}}{R_{t,t+s}} P_{t+s}}_{\text{PV of operating costs}}. \tag{13}$$

Solving for ϵ_t, we obtain:

$$\epsilon_t = \left[\frac{\eta}{1-\eta} \sum_{s=1}^{\infty} \frac{(1-\delta)^{s-1}}{R_{t,t+s}} P_{t+s}\right]^{-\eta}. \tag{14}$$

Hence, the choice of ϵ_t balances the higher initial cost of more efficient capital against the present value of its future energy savings. When the present value of future energy costs is large—because interest rates are low, energy prices are high, or capital is durable—it pays to invest in more efficient capital, and ϵ_t is low.

Equation 14 implies that the elasticity of the energy intensity with respect to the present value of energy costs is η (which governs the trade-off between capital cost and energy efficiency along the technological frontier G; see Equation 4). The elasticity with respect to a transitory one-period interest rate shock—holding compound rates $R_{t,t+s}$ for $s \geq 2$ fixed—is

$$\frac{\partial \log \epsilon_t}{\partial \log R_{t,t+1}} = \eta \frac{P_{t+1}/R_{t,t+1}}{\sum_{s=1}^{\infty} \frac{(1-\delta)^{s-1}}{R_{t,t+s}} P_{t+s}} > 0. \tag{15}$$

Hence, a transitory one-period increase in the interest rate raises the endogenous energy intensity of new investment. Intuitively, the higher interest rate reduces the present value of future energy costs, making it optimal to invest in cheaper, less efficient capital. As a result, the composition of aggregate investment shifts toward a higher energy intensity. This prediction is the central testable implication of the model, which we evaluate empirically in Section 5.

What happens to the energy intensity of *output*, E_{t+1}/Y_{t+1}? This ratio can be decomposed as $(E_{t+1}/K_{t+1}) \times K_{t+1}^{1-\alpha}$. The first factor—average energy intensity per unit of capital—rises because the composition of capital shifts toward less efficient vintages. The second factor, $K_{t+1}^{1-\alpha}$, falls because the capital stock declines. The net effect depends on the quantitative response of the two components.

We note that the effect of interest rate changes on ϵ_t also depends on energy prices, through the interaction of $R_{t,t+1}$ and P_{t+1} in the present-value condition in Equation 13. When energy prices are high, a given interest rate increase has a larger effect on ϵ_t because the discounted energy savings foregone are larger in absolute terms.[6]

Optimal investment and the user cost of capital. The optimal investment equalizes the marginal product of capital (MPK) to an extended measure of the user cost of capital that incorporates the persistent effect of current energy efficiency choices on the future capital stock.

From Equation 7 and Equation 9, using $\chi_t = q_t\lambda_t + \epsilon_t\xi_t$ and shifting time indices, we obtain

$$\overbrace{\alpha K_{t+1}^{\alpha-1}}^{\text{MPK}} = \overbrace{\underbrace{R_{t,t+1}q_t - (1-\delta)q_{t+1}}_{\text{"traditional" user cost}} + \underbrace{\epsilon_t P_{t+1}}_{\text{energy cost}} + \underbrace{(\epsilon_t - \epsilon_{t+1})(1-\delta)\frac{\xi_{t+1}}{\lambda_{t+1}}}_{\text{intertemporal energy cost}}}^{\text{User cost of capital}}, \tag{16}$$

where $\xi_{t+1}/\lambda_{t+1} = \sum_{s=1}^{\infty}(1-\delta)^{s-1}P_{t+1+s}/R_{t+1,t+1+s}$ is the present value of energy costs evaluated at $t+1$.

The user cost of capital has three components. The first is the "traditional" component: by investing at time t the planner forgoes $R_{t,t+1}q_t$ in time $t+1$ consumption, but recoups $(1-\delta)q_{t+1}$ in the form of lower future investment. The second component, $\epsilon_t P_{t+1}$, accounts for the fact that each new unit of capital requires ϵ_t units of energy in its first operating period at price P_{t+1}. The third component captures an intertemporal trade-off implied by time-varying energy efficiency. If energy efficiency is expected to improve over time ($\epsilon_t > \epsilon_{t+1}$), then investing at t rather than waiting until $t+1$ commits the planner to a higher energy cost on the nondepreciated portion of today's investment over its remaining lifetime. This term is positive when $\epsilon_t > \epsilon_{t+1}$, raising the effective user cost and reducing investment—a channel through which anticipated future efficiency gains interact with current investment decisions.

2.3. Impulse responses

We now simulate the model's response to a persistent AR(1) shock to the world interest rate. The shock is an innovation of 100 bps to the gross short-term rate in period 2, with persistence $\rho = 0.80$, so that the path decays gradually over time. This process is calibrated so that the model's interest rate response in net terms roughly matches the empirical IRFs

[6] This interaction motivates an augmented empirical specification with oil prices that we report in the appendix.

in Figures 6–8.[7]

Figure 1 shows the resulting impulse responses. The higher interest rate raises the user cost of capital, reducing investment and output. In an open-economy setting without capital adjustment costs, capital adjusts quickly—firms can reduce K_t until its marginal product aligns with the new world rate—so the drop in K_t is steeper than in output. Energy use, E_t, also falls, tracking K_t closely.

A central feature of the model is the *endogenous* adjustment of the energy intensity ϵ_t. At the time of investment, firms choose ϵ_t along the technological frontier $G(\epsilon_t, q_t) = \eta \log q_t + (1-\eta) \log \epsilon_t = 0$, which embodies the trade-off between higher up-front cost q_t and lower ongoing energy requirements ϵ_t. The optimal choice equates the marginal rate of transformation on this frontier to the present value of future energy costs (Equation 13). A higher interest rate reduces this present value, making it optimal to select cheaper but less efficient capital. This pushes ϵ_t up—energy intensity rises endogenously even as total energy use falls.

The response of E_t/K_t is related but *not identical* to the ϵ_t response. While ϵ_t measures the energy intensity of *new* investment, E_t/K_t reflects the average intensity of the *entire installed capital stock*. Because the stock consists of overlapping vintages, the initial effect of higher ϵ_t is muted by the presence of older, more efficient vintages. As these depreciate, the average intensity E_t/K_t rises, producing a hump-shaped response. Once the pre-shock vintages are largely retired, E_t/K_t gradually converges toward the new steady state. This hump shape is a feature we will look for in the empirical impulse responses.

Meanwhile, the ratio K_t/Y_t falls because capital declines more sharply than output ($Y_t = K_t^\alpha$ with $\alpha < 1$ dampens the output response), and E_t/Y_t also declines because the capital-stock decline dominates the rise in average energy intensity, as discussed in the preceding subsection.

In short, the model predicts that an exogenous interest rate increase generates a prolonged contraction in output, capital, and energy use, alongside a persistent rise in both ϵ_t and E_t/K_t through the endogenous technology choice channel. These model-predicted impulse responses are compared with their empirical counterparts in Section 5.

Remark on directed technical change. In the model, capital vintages are chosen from a fixed menu of technologies. Here we extend the model to allow for dynamic trade-offs with respect to directed technical change. Specifically, we follow Hassler et al. (2021) and

[7] The simulation uses parameters $\alpha = 0.3$, $\beta = 0.96$, $\delta = 0.1$, $\eta = 0.5$, $R_0 = 1.05$, and $P_0 = 1$. The steady state is computed explicitly from the model's optimality conditions rather than via burn-in, ensuring internal consistency between capital, energy use, and the user cost of capital. The PV truncation horizon is $S = 100$ periods, sufficient to capture the effect of future energy costs in the choice of ϵ_t.

Figure 1: *Model impulse responses to a persistent AR(1) interest rate shock.*

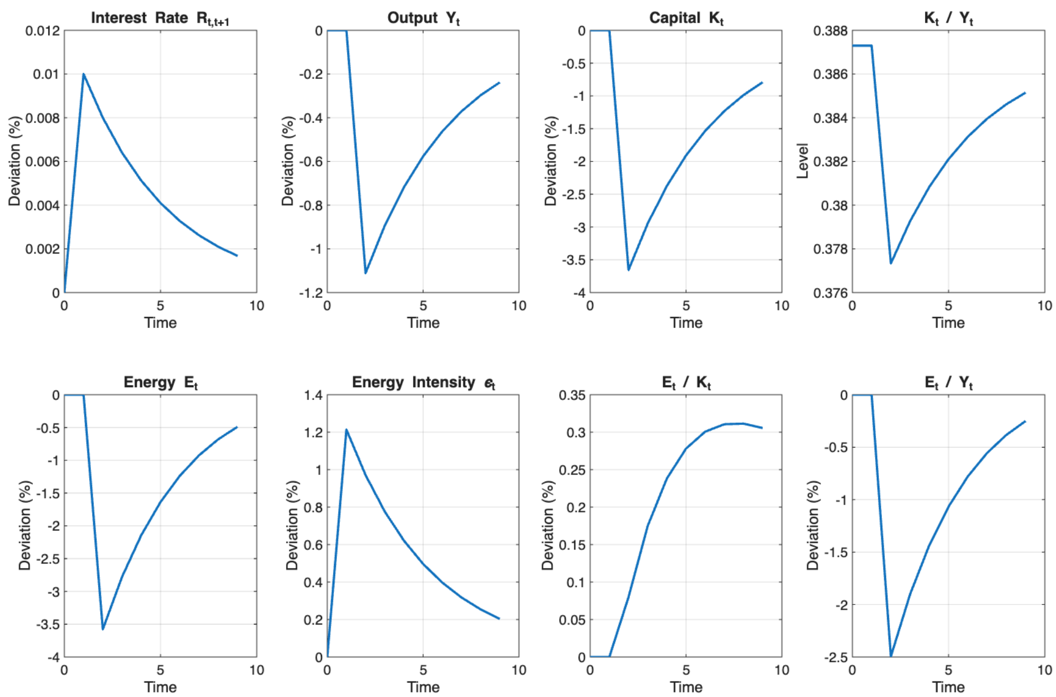

Notes: The interest rate follows an AR(1) process with a 100 bps innovation in period 2 and persistence $\rho = 0.80$, calibrated to match the empirical short-rate responses in Figures **??** and **??**. Panels show the short-term interest rate, output, capital, energy use, energy intensity ϵ_t, and E_t/K_t. Output and capital fall sharply, with rapid capital adjustment in the open-economy setting. Although energy use declines, ϵ_t rises endogenously because higher interest rates tilt investment toward cheaper, less efficient vintages along the $G(\epsilon, q) = 0$ frontier in Equation 13. The E_t/K_t response is hump-shaped due to the gradual retirement of older, more efficient vintages.

assume that the trade-off takes place in the *growth rates* of ϵ_t and q_t rather than their levels:

$$G^{DTC}(\epsilon_t/\epsilon_{t-1}, q_t/q_{t-1}) = \eta \log(q_t/q_{t-1}) + (1-\eta)\log(\epsilon_t/\epsilon_{t-1}) = 0,$$

where G^{DTC} is the technological frontier with directed technical change. In this reformulation, the planner can choose to increase the energy efficiency of new investment (lower $\epsilon_t/\epsilon_{t-1}$) at the cost of increased investment costs (higher q_t/q_{t-1}). The main difference with the baseline is that now the planner's decisions at time t enter as arguments at time $t+1$.

The frontier can be rewritten as

$$\eta \log q_t + (1-\eta)\log \epsilon_t = \eta \log q_{t-1} + (1-\eta)\log \epsilon_{t-1} = \cdots = \eta \log q_{-1} + (1-\eta)\log \epsilon_{-1},$$

where the last equality results from iterating on the first equation until $t = -1$. The technological frontier is therefore static, as in the baseline specification. That is, even though the technology choices embody the possibility of technical change, the model is observationally equivalent to one in which the technological frontier is static.

The finding results from the log-linear functional form for G^{DTC} and is also found in Hassler et al. (2016). The technology frontier implies that any growth in energy efficiency is matched by a proportional increase in the cost of new capital. On the one hand, as energy efficiency grows in one period, it raises the baseline for further growth, making it easier for the planner to choose a more efficient technology. On the other hand, the concavity of the technology frontier means that, in the next period, the largest potential gains are in reducing the cost of capital, and energy efficiency cannot improve as fast. Given the logarithmic form, these two effects cancel exactly, implying a static technological frontier.

In contrast, if G were a linear combination of q_t/q_{t-1} and $\epsilon_t/\epsilon_{t-1}$, the first effect would dominate, and the model could feature path dependency as reductions in energy intensity reduce the cost of further reductions. Alternatively, if the elasticity of substitution between ε_t and q_t were smaller than one, the second effect would dominate. Given sparse evidence on the exact elasticity of substitution, we take the static trade-off as our baseline. For the empirical analysis, the observational equivalence result implies that the model's predictions are robust to whether the underlying technology frontier operates on levels or growth rates, provided the log-linear functional form is a reasonable approximation. As we will see, we do not find evidence for path-dependency, validating our modeling choice.

3. Experimental design

3.1. Data

As articulated in the introduction, we are interested in characterizing how financial conditions affect energy consumption and efficiency, not just in the short run, or at business cycle frequency, but rather extending the analysis further into the medium and longer runs. For this, we rely on long historical annual time series of macroeconomic aggregates across 17 advanced economies.[8] We will use data for macroeconomic and financial variables from the Jordà et al. (2017) database. We refer the reader to the database website for a complete list of variables and sources and will denote the data in this paper as the "JST database" in what follows. The sample reaches back to 1870 for several but not all countries, and

[8] The list of countries is: Australia, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, the United Kingdom, and the United States. Although we have data for Ireland as well, some of the CO2 emissions data is missing.

we stop the sample in 2019 to remove the effects of the COVID-19 pandemic. We will also exclude from the analysis the world war periods (1914-1918 and 1939-1945).

In addition, we merge data from the long-term productivity database of Bergeaud et al. (2016), maintained by the Banque de France. The database covers essentially the same countries as the JST database since 1890 at annual frequency and contains observations on total factor productivity per hour worked, labor productivity per hour worked, capital intensity and GDP per capita. Usefully, they also have data on capital broken down by "machines" and "buildings." The extent to which technological advances in energy use typically center around "machines" will provide a convenient cross-check of our model's intuition as we will show.

We collect data on energy consumption from Our World in Data, which consolidates data from the U.S. Energy Information Administration - International Energy Data (2023) and the Energy Institute - Statistical Review of World Energy (2024) for energy use. Because data on primary energy consumption have only been available since 1965, we rely on the same sources to obtain data for CO2 emissions, which cover the same sample as the JST database.[9] As we discuss below, CO2 emissions correlate well with energy use, which allows us to better explore the time sample properties of our macro series.

3.2. Emissions and energy consumption

As a first glimpse at the evolution of CO2 emissions (which we shall use as a proxy for energy consumption before 1965), Figure 2 showcases the 17-country aggregate total CO2 emissions expressed as a ratio to real GDP measured in PPP-adjusted 1990 dollars. The figure shows an index based on this ratio that is normalized to 100 in 1965, which is the first year data on energy use are available (so as to make comparisons across figures easier).

This figure illustrates several features of the evolution of CO2 emissions in these large economies. First, note that, on a per-dollar-produced basis, emissions in 2020 are about two thirds lower than they were in 1965, likely reflecting the gradual shift in the contribution of services in the economy relative to farming and manufacturing, among other factors that probably include gains in energy efficiency. Appendix A1 shows the same figure broken down by country in Figures A.1–A.4. Many feature the same bell-shaped pattern of Figure 2. The case of the U.K. is interesting since it was the first economy to enter the industrial revolution. Emissions intensity has been gradually declining throughout the entire sample rather than being bell-shaped. In the U.S., emissions peaked around the end of World War I as the country transitioned from coal to oil, whereas in some of the

[9] OWID also relies on Andrew and Peters (2024) for data on CO2 emissions.

Figure 2: *Emissions intensity measured as CO2 emissions over real GDP*

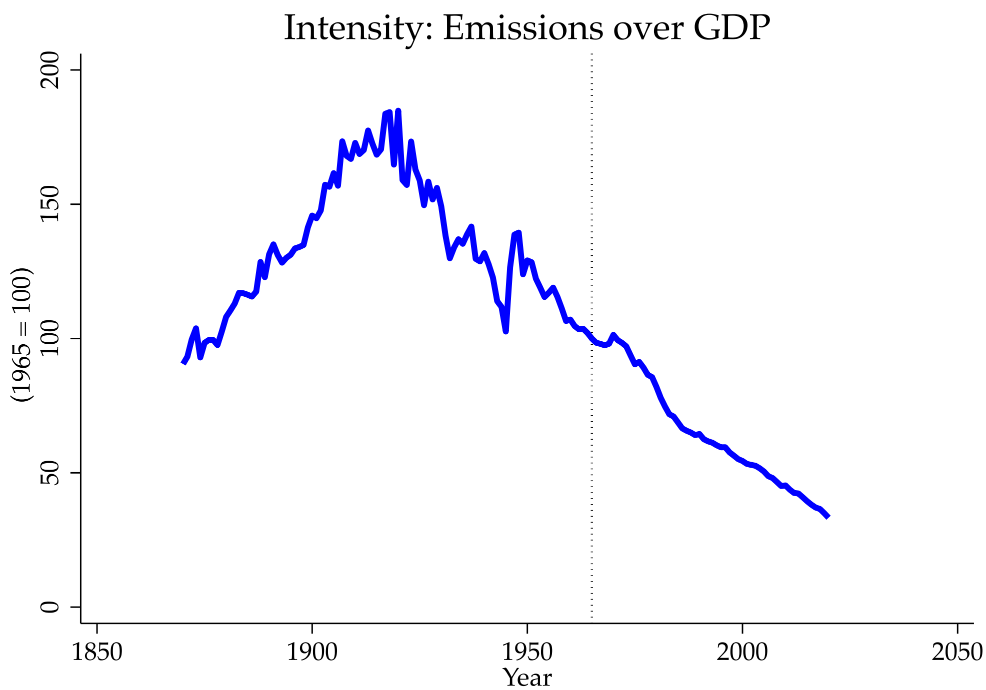

Notes: Total CO2 emissions for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. Data normalized to 100 in 1965. See text.

economies that developed later (*e.g.*, Finland, Italy, or Spain) the peak of the bell curve happens well after World War II.

However, our paper focuses on energy efficiency not on CO2 emissions directly. Because of data limitations, primary energy consumption measured in terawatt-hours is only available since 1965. However, as one would expect, there is a tight connection between CO2 emissions and efficiency. We illustrate this connection in Figure 3, which displays total CO2 emissions and primary energy consumption, both normalized to 100 in 1965.

The figure shows that, to a first approximation, one can use CO2 emissions as a proxy for energy consumption though as Appendix A2 shows, there can be considerable variation across countries. Appendix A2 displays country-specific figures of CO2 emissions versus energy consumption relative to real GDP in Figures A.5–A.8. These figures show that, whereas for some countries the time series of CO2 emissions and energy consumption track closely, for others there is more of a difference in their evolution.

Given these considerations, in what follows we will mostly focus our analysis on results based on CO2 emissions to take advantage of its better data coverage. At times, however, we will also compare those findings to results using energy consumption.

Figure 3: *CO2 emissions versus primary energy consumption*

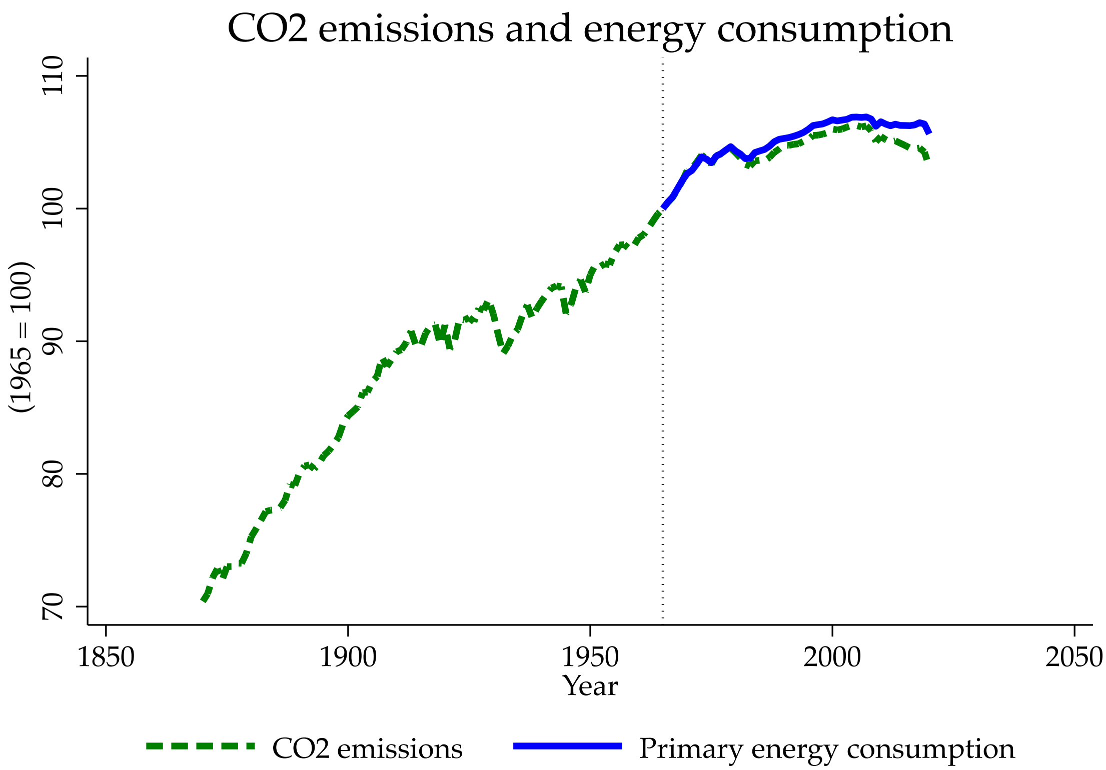

Notes: Total CO2 emissions and primary energy consumption for the 17 countries in the JST database. Data normalized to 100 in 1965. See text.

3.3. CO2 emissions over the business cycle

As a first pass only meant to motivate further analysis later on, we report a simple exercise to get a better sense of the business cycle dynamics of CO2 emissions. Using the full sample of countries and years available, we focus on the average path of an economy from the onset of a recession to 8 years since its start. We plot the paths of CO2 emissions, real GDP, and investment to GDP ratio in Figure 4. The figure shows the cross-country average paths followed by these three variables, where time zero marks the start of a recession. The latter is set to the first year in which real GDP growth turns negative.

Several features stand out. First, the initial GDP decline is in the order of 3 percent but the recovery comes, on average, swiftly. By year 3 the economy is back to the level it started with before the recession, and by year 8, it is about 12 percent higher. The investment to GDP ratio, however, suffers a more dramatic and longer lasting downturn, with a recovery to pre-recession levels only starting by year 7. CO2 emissions track the path of real GDP quite closely initially, however, by year 8, emissions start to diverge relative to the growth

Figure 4: *Business cycle dynamics of CO2 emissions*

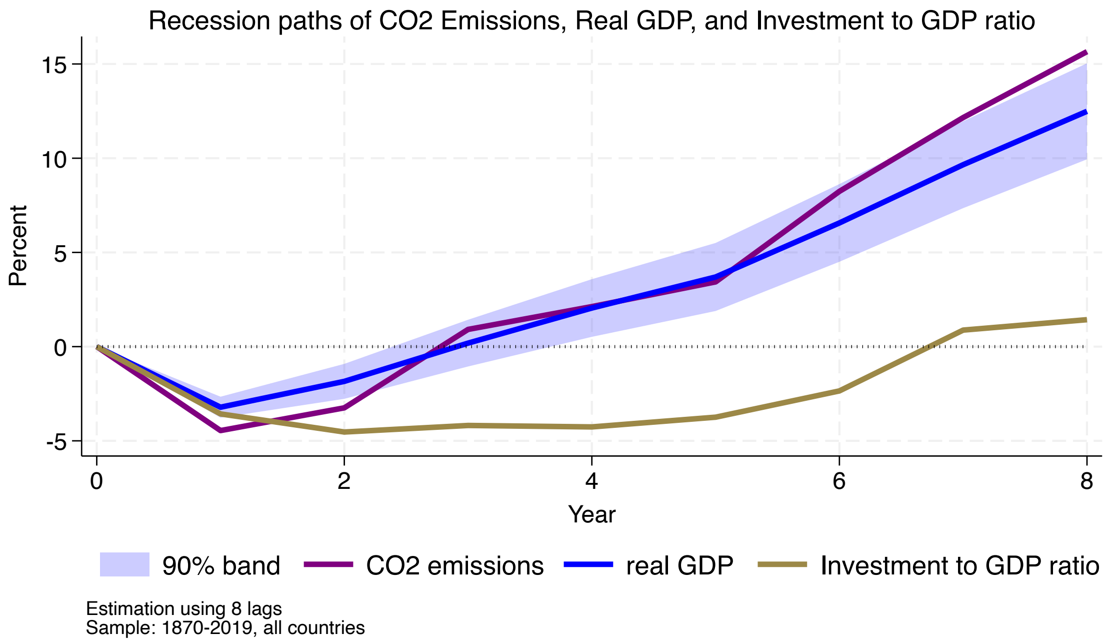

Notes: The figure displays the percentage change in total CO2 emissions, real GDP and the investment to GDP ratio since the start of a recession. See text.

in the economy.

Recessions illustrate that in times of economic hardship, investment suffers by an even greater amount over a longer period of time. We may thus conjecture that this retrenchment in investment may have consequences on the type of capital investment, energy use and eventually, CO2 emissions. Note that this should be understood to be relative to the trends that we displayed in Figure 2.

3.4. Monetary stance

For our empirical analysis, we focus on changes in financial conditions induced by exogenous variation in the stance of monetary policy. There is no universally accepted way to measure the monetary stance, that is, whether monetary policy is *expansionary*, *contractionary*, or *neutral*. However, there is some consensus that fluctuations of interest rates around some measure of the *natural rate of interest* seem like a sensible place to start. To quote Holston et al. (2017) "[t]he natural or equilibrium real interest rate provides a benchmark for measuring the stance of monetary policy, with policy expansionary (contractionary) if the short-term real interest rate lies below (above) the natural rate" (p. S59). Rather than developing our own measures of the natural rate and the stance variable, we borrow directly from Grimm et al. (2023), which in turn expand measures of the natural rate from Del Negro et al. (2019).

More specifically, Del Negro et al. (2019) identify country-specific and global trends in interest rates and inflation, and their corresponding stationary components, in a VAR model with common trends. This is done by imposing long-run restrictions on the short and long ends of the yield curve. This trend-cycle decomposition is performed by Bayesian estimation. Grimm et al. (2023) extend the original analysis from 7 to 18 advanced economies (the same set of countries we use in our dataset). Since we borrow directly from their analysis, we refer the reader to the original sources for more details. Here we only give a brief synopsis of the approach.

Denote $r^*_{i,t}$ the real natural rate of interest for country i at time t obtained from Grimm et al. (2023). Next, we use their definition of stance, which we repeat here for convenience:

$$s_{i,t} = \overline{stance}_{i,t} = \frac{1}{5}\sum_{k=0}^{4}\left(r_{i,t-k} - r^*_{i,t-k}\right). \tag{17}$$

The reason for taking a 5-year average (recall that our data are annual), is to smooth out short-run fluctuations in the stance and instead focus on protracted periods where the stance is tight/loose. As Grimm et al. (2023) report, the results are relatively robust to choosing an average over a somewhat shorter or longer time span. Figure 5 replicates Figure 2 in Grimm et al. (2023), which displays the Del Negro et al. (2019) estimates of r^* along with the Grimm et al. (2023) estimates that we use in this paper.

4. Statistical methods

We estimate the dynamic multipliers of the main variables of interest using local projections Jordà (2005). Let $y_{i,t}$ denote an outcome variable of interest, for example, 100 times the log of real GDP in country i in year t. We transform this outcome variable by taking the *long difference*, that is $\Delta_h y_{i,t+h} \equiv y_{i,t+h} - y_{i,t-1}$ for $h = 0, 1, \ldots, H$. This transformation, in our example, would represent the approximate percentage change in real GDP from $t-1$ to $t+h$. Using long-differences has two consequences. First, it generates smoother responses since this transformation can be seen as a form of averaging across yearly changes (Jordà and Taylor, 2024). Second, Piger and Stockwell (2023) show that long-differencing greatly attenuates small sample biases when the data are persistent.

Next, we define the *treatment* variable as $s_{i,t}$, in our case, a shock to the stance variable defined in Equation 17. Finally, we collect all other conditioning information in the vector $\Delta \boldsymbol{x}_{i,t}$. This vector contains lags of the outcome, the treatment, and all other exogenous

Figure 5: *Comparing the r-star estimates from Del Negro et al. (2019) and Grimm et al. (2023)*

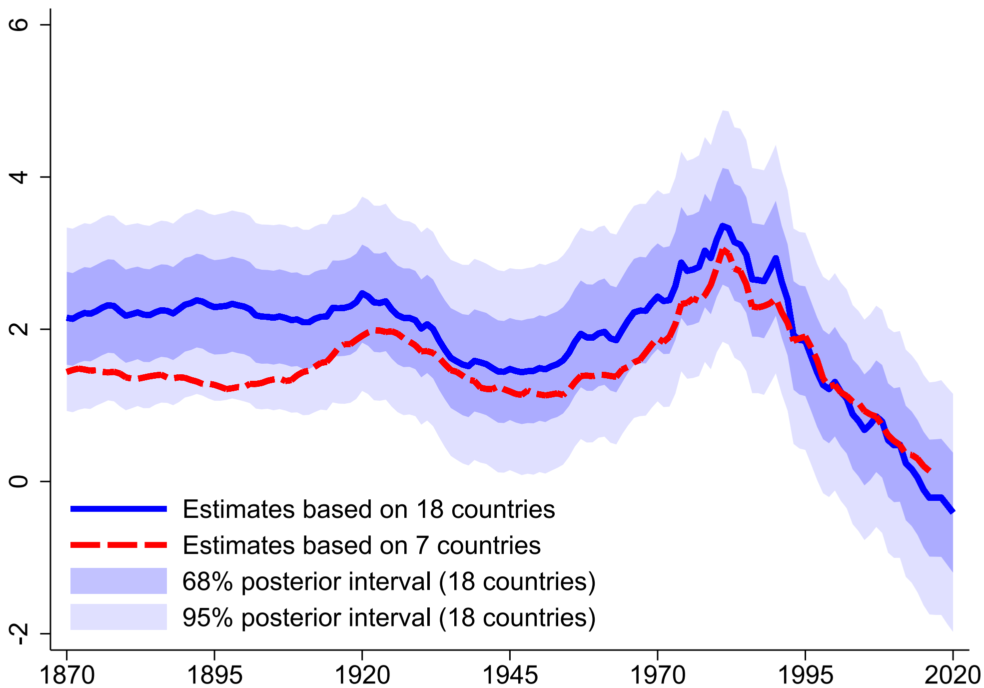

Notes: This figure replicates Figure 2 in Grimm et al. (2023) and displays r-star estimates based on a sample of 7 advanced economies using the Del Negro et al. (2019) original sample versus the 18 advanced economies used in Grimm et al. (2023). See text.

and predetermined variables. In our specification, we also allow for country fixed effects. In order to control for time fixed effects and to economize on parameters, we include a measure of *global* real GDP growth (computed by adding the GDP of all countries in our sample), denoted as Δg_t.

Given these definitions, the typical specification of the local projection is:

$$\Delta_h y_{i,t+h} = \alpha_{i,h} + \delta_h \Delta g_t + \beta_h s_{i,t} + \gamma_h' \Delta \boldsymbol{x}_{i,t} + u_{i,t+h} \tag{18}$$

where β_h is the parameter of interest for the dynamic multiplier of cumulative impulse response:

$$\mathcal{R}_{s\to y}(h) \equiv \mathbb{E}[\Delta_h y_{i,t+h} | s_{i,t} = 1; \Delta \boldsymbol{x}_{i,t}] - \mathbb{E}[\Delta_h y_{i,t+h} | s_{i,t} = 0; \Delta \boldsymbol{x}_{i,t}] = \beta_h. \tag{19}$$

Of course, we expect that, despite including a rich set of controls (to be specified momentarily), $s_{i,t}$ might not meet a selection-on-observables condition, either because of omitted unobservables or due to simultaneity, for example. Thus, to address endogenenity concerns, we estimate Equation 18 using instrumental variables methods, as the next section presents in detail.

The model presented in Section 2 also predicts that the effect of tight monetary conditions is amplified whenever the prices of energy are high. Though not central to our analysis, we note here that we will extend the specification in Equation 18 as follows:

$$\Delta_h y_{i,t+h} = \alpha_{i,h} + \delta_h \Delta g_t + \beta_h s_{i,t} + \phi_h s_{i,t}(\Delta p_t^{oil} - \Delta \overline{p}^{oil}) + \gamma_h' \Delta \boldsymbol{x}_{i,t} + u_{i,t+h}, \tag{20}$$

where p_t^{oil} refers to the price of crude oil, measured in constant 2023 US$ per cubic meter (a barrel of oil is approximately 0.159 cubic meters). The data for crude oil come from the Energy Institute based on S&P Global Platts - Statistical Review of World Energy (2025) with major processing from Our World in Data.[10]

We experiment with a setting where oil prices are 50% above their sample mean, which turns out to be approximately two standard deviations. Thus, the equivalent to Equation 19 in this situation can be calculated by noting that:

$$\mathcal{R}_{s\to y}(h; \Delta p_t^{oil} = \Delta \overline{p}^{oil} + 50) = \beta_h + \phi_h 50. \tag{21}$$

We conclude this section by listing the variables to be included in $\Delta \boldsymbol{x}_{i,t}$. These are: 100 times the logs of real GDP, the consumer price index (CPI), total CO2 emissions, total private credit to the nonfinancial sector, and capital. In addition, we include the 100 times the change in the investment to GDP ratio, and the change in the short-term interest rate. We include up to 8 lags of these variables. This is the length favored by the AIC criterion in 7 out of 17 countries. Moreover, since 8 is the length of the impulse response that we plot, results in Lewis and Reinsel (1985) and Jordà et al. (2024) suggest that even when the lag length is truncated, consistency is preserved for horizons equal to or smaller than truncation lag.

4.1. Constructing the instrument

What might be a good source of exogenous fluctuations in the *stance* variable, $s_{i,t}$? Obstfeld and Taylor (2004) provide a possible answer through the concept of the *trilemma* of international finance: countries cannot simultaneously fix their exchange rate, allow for unfettered movement of capital across borders, and maintain control over domestic interest rates.

[10] These data can be downloaded from:
https://ourworldindata.org/grapher/oil-prices-inflation-adjusted.

The logic of the trilemma is disarmingly simple. With fixed exchange rates and free movement of capital, any differences in the returns of assets of similar risk characteristics will be arbitraged away by borrowing assets with low returns and investing in assets with high returns. Jordà et al. (2020) used this idea to construct an instrument for fluctuations in interest rates and Grimm et al. (2023) refined the approach further. Here, we sketch the main ideas and refer the reader to these papers for greater detail.

Let $f_{i,t} \in \{0, 1\}$ indicate whether a country pegs its exchange rate in years t and $t-1$ (to minimize instances of opportunistic pegging). Let $o_{i,t} \in [0, 1]$ indicate the degree of capital openness based on a rescaled version of the index provided in Quinn et al. (2011).[11] Let $R_{b(i,t),t}$ denote the nominal interest rate of the base economy that country i uses to peg its exchange rate to at time t. Let $\hat{R}_{b(i,t),t}$ denote the predicted value of $R_{b(i,t),t}$ using base-country domestic variables. Think of it as fitting an extended Taylor rule. The idea is that $\hat{\delta}_{b(i,t),t} = R_{b(i,t),t} - \hat{R}_{b(i,t),t}$ represents the component of the base-country interest rate fluctuations that could not be predicted by domestic considerations (such as base-country inflation deviations from target, output growth, and so on). A complete list of predictors is provided in Grimm et al. (2023) and include lags of output growth, lags of inflation, lags of interest rates, etc.

With these definitions, the instrument is defined as:

$$z_{i,t} = f_{i,t} \times o_{i,t} \times \hat{\delta}_{b(i,t),t}. \tag{22}$$

Specifically, we use the unpredictable component of the base-country interest rate fluctuations, scaled by the degree of capital mobility, whenever a country i pegs to a given base country b. A table with the definition of base and pegging economies throughout our sample is provided in Jordà et al. (2020). First stage tests on the strength of the instrument indicate that the instrument is quite strong (the specifics can be found in Grimm et al., 2023). For example, the first stage regression of the stance variable on the instrument set used in several of our regressions (using panel data and clustered robust standard errors) yields an F-statistic of 22.02. Moreover, Grimm et al. (2023) show that the pass-through implied by their first-stage regressions is consistent with the economic logic of the trilemma and results reported in previous research.

[11]This index does not include Ireland, which heretofore is excluded in the construction on the instrumental variable.

5. Empirical results

5.1. Full sample

We begin with the full sample (1890–2019, excluding World Wars) for the 17 economies described in Section 3.1. We report here estimates of the responses based on Equation 18 and leave estimates that augment the original specification as described in Equation 20 with oil prices to the appendix.

Figure 6 displays cumulative impulse responses to a 1 percentage point increase in the monetary stance for the main aggregates. The figure shows the responses of levels—the real short-term interest rate, real GDP, real capital, and CO2 emissions—as well as the intensity ratios central to our analysis: K/GDP, CO2/K, and CO2/GDP.[12]

The level responses accord with standard macroeconomic intuition. The real short-term interest rate rises (shown in panel (a) of Figure 6) by about 2 percentage points, peaking around year 3–4 before gradually reverting. Real GDP (shown in panel (b)) declines cumulatively by approximately 6% over 8 years. Real capital (panel (c)) falls by roughly 10%—nearly twice the GDP decline—reflecting the direct sensitivity of investment to the user cost of capital. CO2 emissions (shown in panel (d)) decline by about 10%, tracking capital more closely than output.

The intensity ratios in panels (e)–(g) of Figure 6 deliver the key results. The decline in K/GDP (panel (e)) confirms that capital contracts proportionally more than output. Most importantly, the CO2/K ratio (panel (f)) *rises* by approximately 3% by year 3, indicating that the composition of the capital stock shifts toward more emissions-intensive vintages—precisely the vintage-choice mechanism formalized in Section 2. The CO2/GDP ratio (panel (g)) exhibits a more muted response, as the effects of declining output and rising emissions intensity partially offset one another.

In all panels, shaded regions denote 68% confidence bands ($\pm$ one standard error) and dashed lines mark 90% confidence bands. These are pointwise and reflect estimation uncertainty at each individual horizon; they do not permit joint statements about the trajectory. A one-standard-deviation region contains approximately two thirds of the probability mass, making it a useful gauge of economic significance beyond statistical significance. Results incorporating oil price interactions, reported in Figure A.9, show broadly similar patterns with somewhat larger capital declines and a slightly tighter interest rate path when oil prices are elevated.

[12] We calculate the real interest rate by using an AR(8) model to predict inflation and hence define the real rate as $r_t = i_t - E_t\pi_{t+1}$, where r_t is the real rate, i_t is the nominal rate, and $E_t\pi_{t+1}$ is the one-period ahead inflation forecast from the AR(8) model.

Figure 6: *Cumulative response of main aggregates to a monetary stance shock.*

Full sample: 1890–2019

(a) *Short-term Real Interest Rate*

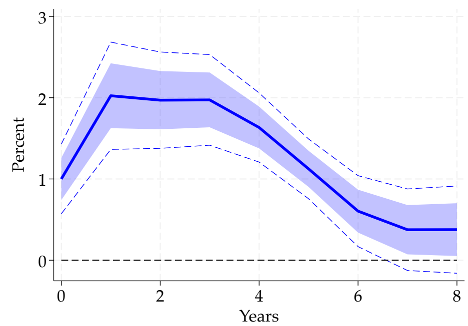

(b) *Real GDP*

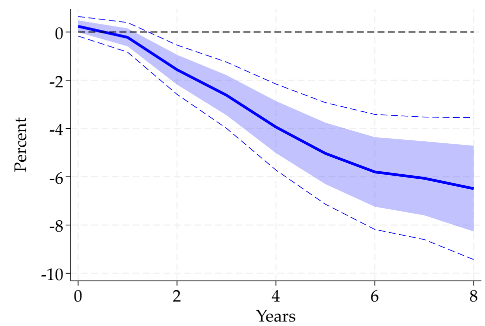

(c) *Real Capital (K)*

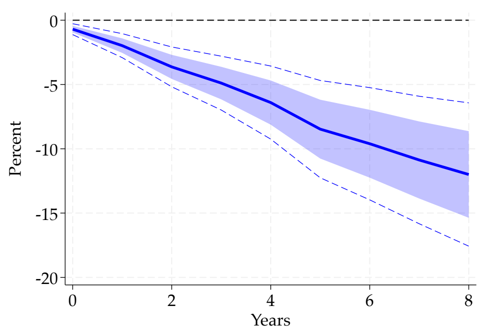

(d) *CO2 Emissions*

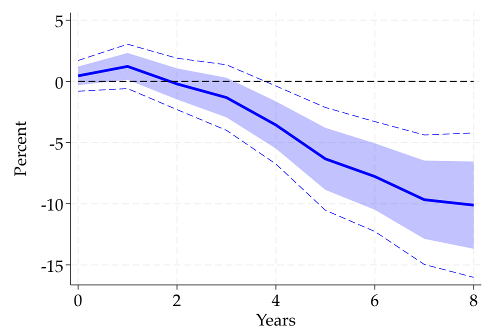

(e) *Capital to GDP ratio*

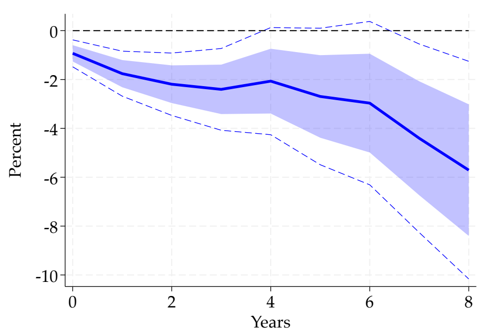

(f) *CO2 Emissions to Capital ratio*

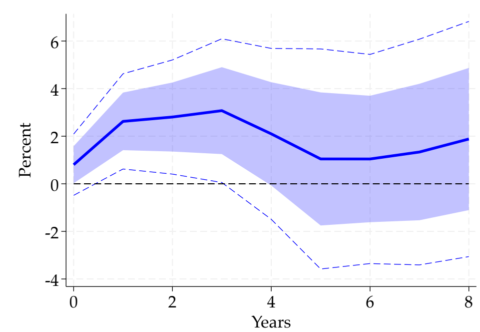

(g) *CO2 Emissions to GDP ratio*

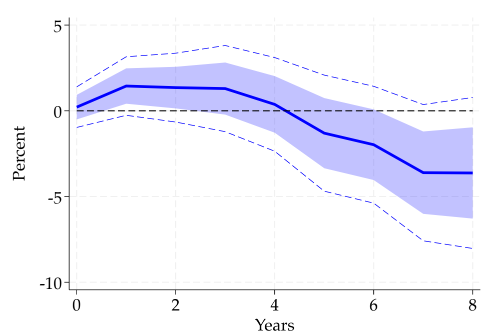

Notes: Cumulative responses to a shock in stance instrumented with the trilemma instrument. Shaded area: 68% CI (± one standard error). Dashed lines: 90% CI. Sample: 1890–2019 excluding World Wars. See text.

5.2. Post-1965 sample

Next, we examine the sample that directly contains data on energy consumption, from 1965 to 2019. This allows us to verify that indeed the responses based on CO2 emissions are a good proxy for energy consumption and also to investigate a more recent period, where financial market developments and monetary policy are conducted more similarly to current times. This will also serve as an informal check for potential structural breaks.

Figure 7 displays cumulative responses of the main aggregates, that is, the short-term interest rate, real GDP, the real capital stock, CO2 emissions and total energy. Figure 8 displays the relevant variables explored in our theoretical model, namely the following ratios: capital to output, CO2 emissions to capital, CO2 emissions to GDP, energy to capital, and energy to output.

Qualitatively, the level responses mirror those of the full sample. Quantitatively, the effects are attenuated: real GDP (shown in panel (b) of Figure 7) declines by about 2% over 8 years (versus 6%), and real capital (panel (c)) falls by roughly 4% (versus 10%). CO2 emissions (panel (d)) decline by about 3%, closely tracking the GDP response. Total energy consumption (panel (e)) follows a nearly identical path to CO2 emissions, confirming the tight link between the two documented in Figure 3 and validating the use of CO2 as a proxy for energy in the full sample. Figure A.11 shows that when oil prices are 50% above their mean, the capital decline is closer to 6%.

Figure 8 the results most directly linked to the model. The E/K ratio—the empirical counterpart to E_t/K_t in Section 2 and shown in panel (d) of Figure 8—rises by approximately 3% around year 6, exhibiting the delayed, hump-shaped profile predicted by the vintage capital mechanism. The CO2/K ratio (panel (b)) displays a similar pattern. Meanwhile, E/GDP (panel (e)) and CO2/GDP (panel (c)) remain close to zero or slightly negative, indicating that the composition effect on energy intensity operates primarily through the capital channel. This accords with the model's prediction that both the numerator and denominator of E_t/Y_t decline, with the concavity of the production function ($\alpha < 1$) causing the denominator to fall proportionally less than capital.

The attenuation of level responses between the full and post-1965 samples likely reflects structural changes in the monetary transmission mechanism as well as reduced statistical power from the shorter sample. Critically, however, the qualitative patterns—and in particular the sign, shape, and approximate magnitude of the E/K and CO2/K responses—are robust across both samples.

Figure 7: *Cumulative response of main aggregates to a monetary stance shock.*

Post-1965 sample: 1965–2019

(a) *Short-term interest rate*

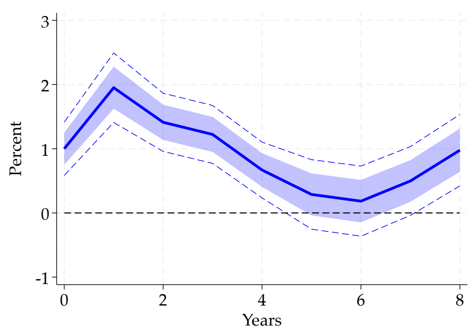

(b) *Real GDP*

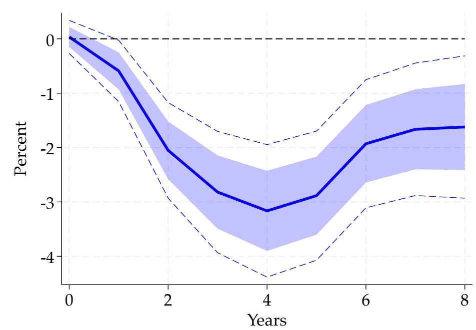

(c) *Real Capital*

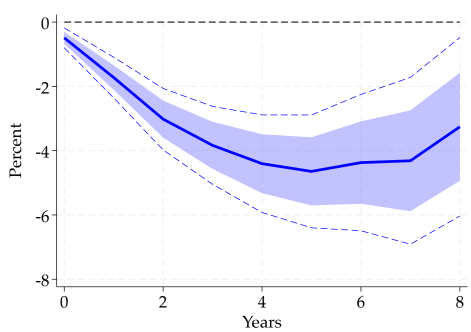

(d) *CO2 emissions*

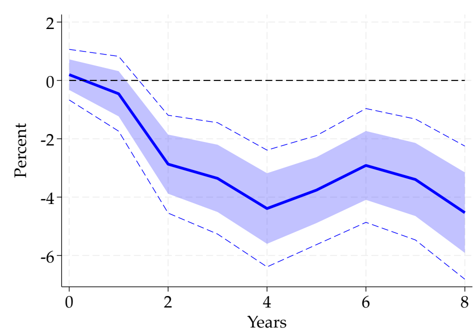

(e) *Total energy*

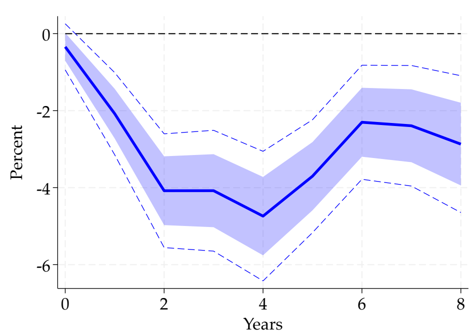

Notes: Cumulative responses to a shock in stance instrumented with the trilemma instrument. Shaded area: 68% CI (± one standard error). Dashed lines: 90% CI. Sample: 1965–2019. See text.

Figure 8: *Cumulative response of main aggregates to a monetary stance shock.*

Post-1965 sample: 1965–2019

(a) *Capital to GDP ratio*

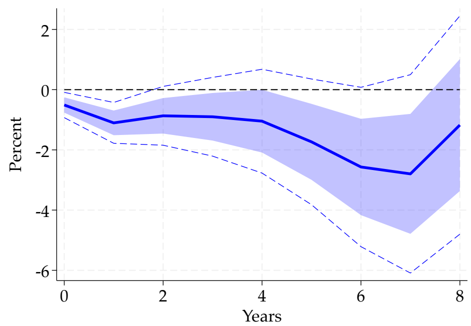

(b) *Emissions to GDP ratio*

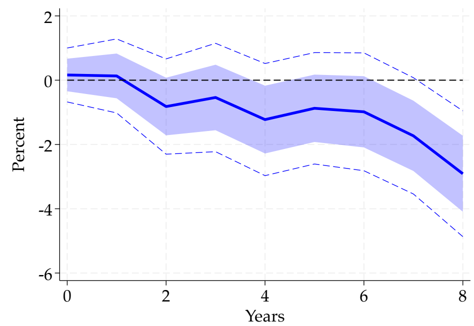

(c) *Emissions to GDP ratio*

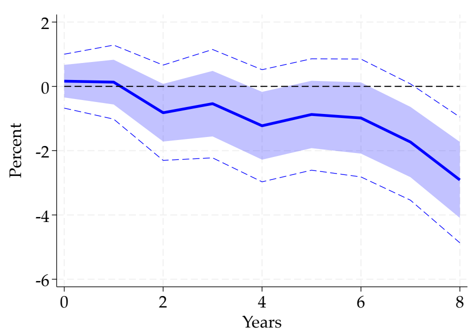

(d) *Energy to capital ratio*

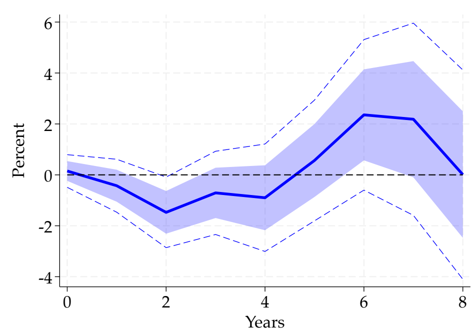

(e) *Energy to GDP ratio*

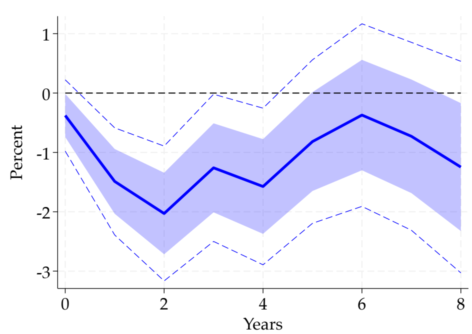

Notes: Cumulative responses to a shock in stance instrumented with the trilemma instrument. Shaded area: 68% CI (± one standard error). Dashed lines: 90% CI. Sample: 1965–2019. See text.

5.3. Mapping to model

The model in Section 2 is intentionally parsimonious: a small open economy with vintage capital, no adjustment costs, and a simple AR(1) interest rate process. The goal is not to generate a quantitative match by adding numerous frictions, but rather to isolate the mechanism linking financial conditions to energy efficiency choices. We now compare the theoretical impulse responses in Figure 1 with the empirical evidence in Figures 6 and 8, organized around what the model does and does not capture.

The model reproduces the main empirical patterns well. First, a persistent monetary tightening generates sustained declines in output, capital, and energy use, consistent with both empirical samples. Second, capital contracts proportionally more than output, so that K/Y falls—reflecting the direct sensitivity of investment to the user cost of capital in both model and data. Third, energy and emissions track capital rather than output, declining by magnitudes closer to the capital response.

Fourth, and most importantly for the paper's central message, the model's E_t/K_t ratio rises with a hump-shaped profile and a delayed peak—matching the empirical pattern of the E/K and CO2/K responses. In the model, this hump arises because newly installed vintages carry higher energy intensity ϵ_t (the optimal ϵ_t rises when higher interest rates reduce the present value of future energy savings, as shown in Equation 14), but their effect on the *aggregate* ratio E_t/K_t builds only gradually as older, more efficient vintages depreciate. Also consistent with the model, E/K and CO2/K responses converge back to trend as conditions normalize, with no indication of permanent effects or path dependency. The delayed empirical peak around year 6 aligns with this vintage-composition mechanism. Fifth, the modest response of E/GDP and CO2/GDP in the data is consistent with the model's E_t/Y_t panel, where both energy and output decline but the concavity of the production function attenuates the net effect on their ratio.

The model underestimates all magnitudes, as expected from its stripped-down structure. The empirical GDP declines of 2–6% (depending on the sample) and capital declines of 4–10% substantially exceed the model's approximately 1% output and 3.5% capital declines. The model's E_t/K_t increase peaks at about 0.3 percentage points, roughly an order of magnitude below the empirical peaks of approximately 3%.

Three deliberate simplifications in the model account for most of these gaps. First, and most directly, the AR(1) interest rate process generates a monotonically decaying rate path, whereas the empirical interest rate response is clearly hump-shaped, peaking around year 3–4 before reverting. Since the interest rate path drives all downstream responses through the present-value condition for ϵ_t and the user cost of capital, a more persistent tightening would amplify both the level and the composition effects. Second, the absence of capital adjustment costs allows capital to jump immediately to equate the marginal product to the world interest rate. In practice, adjustment frictions slow capital reallocation and generate the more gradual, drawn-out responses observed in the data. Third, the model abstracts from variable capital utilization. In a downturn, firms would plausibly use their most efficient vintages more intensively, which would attenuate the observed rise in average energy intensity—suggesting that the *underlying* vintage-choice effect may be even larger than what E_t/K_t reveals.

Despite these quantitative gaps, the broad qualitative alignment across two independent samples spanning over a century of data provides support for the model's core mechanism: higher interest rates tilt investment toward cheaper but less energy-efficient capital by increasing the discount rate applied to future energy savings. This composition effect generates a persistent rise in energy intensity relative to capital, even as total energy consumption falls. The model's ability to reproduce the direction, shape, and relative timing of the key responses—without any calibration to match the empirical impulse responses—suggests that the vintage-choice channel is an empirically relevant mechanism linking financial conditions to energy efficiency.

6. Conclusion

This paper shows that financial conditions shape not only how much firms invest, but also what they invest in: higher interest rates tilt new capital toward lower up-front cost and higher energy intensity because they shrink the present value of future energy savings. Using 150 years of data for 17 advanced economies, using CO_2 emissions as a proxy for energy consumption, and a trilemma-based instrumental variable identification strategy, we confirm this mechanism. Monetary tightening reduces output, capital, and total energy use, yet raises energy intensity per unit of capital via a composition effect that is persistent and hump-shaped, peaking around 3% about six years after the shock (based on the post-1965 sample results). The timing reflects slow capital turnover, and the effect is stronger for machinery than for buildings.

Our results point to a policy-relevant tradeoff. The model as written abstracts from the emissions externality, but incorporating it is straightforward—for example, by letting the energy price path P_{t+s} include a Pigouvian carbon component. With that extension in mind, two implications follow immediately. First, the short-run emissions reduction from a monetary tightening can overstate the medium-run benefit: while energy use falls with output in the contraction, the capital installed during the episode is less energy-efficient, so the recovery occurs with a worse vintage mix and higher emissions intensity than would otherwise prevail. Second, the model clarifies the interaction between interest rates and carbon pricing. Our model implies that the efficiency choice depends on the present value $\sum_{s=1}^{\infty}(1-\delta)^{s-1}P_{t+s}/R_{t,t+s}$: carbon pricing raises P_{t+s} and strengthens the incentive to adopt efficient capital, whereas higher interest rates raise $R_{t,t+s}$ and attenuate that incentive. Because both enter the same present-value term, carbon pricing is less effective at steering the composition of new investment toward energy efficiency when discount rates are high. This interaction is complementary to policy analyses of the green transition that emphasize diffusion across vintages via investment, scrapping, and innovation incentives (e.g., Keuschnigg and Stalenis 2025); incorporating discount-rate variation into such environments is a natural direction for future work.

References

Aghion, P., Bergeaud, A., De Ridder, M., and Van Reenen, J. (2024). Lost in transition: Financial barriers to green growth. Technical report, INSEAD Working Paper.

Aghion, P., Boneva, L., Breckenfelder, J., Laeven, L., Olovsson, C., Popov, A., and Rancoita, E. (2022). Financial markets and green innovation. *ECB working paper*, (2686).

Aghion, P., Farhi, E., and Kharroubi, E. (2019). Monetary policy, product market competition and growth. *Economica*, 86(343):431–470.

Amador, S. (2022). Hysteresis, endogenous growth, and monetary policy. Technical report, U.C. Davis.

Andrew, R. and Peters, G. (2021). The global carbon project's fossil CO2 emissions dataset: 2021 release. *CICERO Center for International Climate Research*. Internal version GCB2021v32.

Andrew, R. and Peters, G. P. (2024). The global carbon projects fossil co 2 emissions dataset: 2024 release. *CICERO Center for International Climate Research, Oslo*.

Andrew, R. M. (2019). Global CO2 emissions from cement production, 1928–2018. *Earth System Science Data*, 11(4):1675–1710.

Atkeson, A. and Kehoe, P. J. (1999). Models of energy use: Putty-putty versus putty-clay. *American Economic Review*, 89(4):1028–1043.

Bergeaud, A., Cette, G., and Lecat, R. (2016). Productivity trends in advanced countries between 1890 and 2012. *Review of Income and Wealth*, 62(3):420–444.

Bernanke, B. S. and Mihov, I. (1998). The liquidity effect and long-run neutrality. In *Carnegie-Rochester conference series on public policy*, volume 49, pages 149–194. Elsevier.

Bilal, A. and Stock, J. H. (2024). Macroeconomics and climate change. Technical report, National Academies of Sciences, Engineering, and Medicine.

Blanchard, O. J. and Summers, L. H. (1986). Hysteresis and the european unemployment problem. *NBER Macroeconomics Annual*, 1:15–78.

BP (2021). Statistical review of world energy.

Brunnermeier, M., Palia, D., Sastry, K. A., and Sims, C. A. (2021). Feedbacks: financial markets and economic activity. *American Economic Review*, 111(6):1845–1879.

Cakir Melek, N. and Orak, M. (2024). Technology and energy substitution: A path toward climate change mitigation. *Federal Reserve Bank of Kansas City Working Paper*, (23-15).

Casey, G. (2024). Energy efficiency and directed technical change: implications for climate change mitigation. *Review of Economic Studies*, 91(1):192–228.

Del Negro, M., Giannone, D., Giannoni, M. P., and Tambalotti, A. (2019). Global trends in interest rates. *Journal of International Economics*, 118:248–262.

Etemad, B., Luciani, J., Bairoch, P., and Toutain, J.-C. (1991). *World Energy Production 1800-1985*. Droz, Geneva.

Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Garga, V. and Singh, S. R. (2021). Output hysteresis and optimal monetary policy. *Journal of Monetary Economics*, 117:871–886.

Gilfillan, D. and Marland, G. (2021). Cdiac-ff: Global and national CO2 emissions from fossil fuel combustion and cement manufacture: 1751-2017. *Scientific Data*, 8(1):2.

Grimm, M., Jordà, Ò., Schularick, M., and Taylor, A. M. (2023). Loose monetary policy and financial instability. Technical report, National Bureau of Economic Research.

Hassler, J., Krusell, P., and Olovsson, C. (2021). Directed technical change as a response to natural resource scarcity. *Journal of Political Economy*, 129(11):3039–3076.

Hassler, J., Krusell, P., and Smith Jr, A. A. (2016). Environmental macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 1893–2008. Elsevier.

Holston, K., Laubach, T., and Williams, J. C. (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*, 108:S59–S75.

IPCC (2006a). 2006 IPCC guidelines for national greenhouse gas inventories, volume 2, chapter 6: Reference approach. Technical report, Intergovernmental Panel on Climate Change.

IPCC (2006b). 2006 IPCC guidelines for national greenhouse gas inventories, volume 2: Energy. Technical report, Intergovernmental Panel on Climate Change.

Jones, C. I. (2005). The shape of production functions and the direction of technical change. *The Quarterly Journal of Economics*, 120(2):517–549.

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

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

Jordà, Ò., Schularick, M., and Taylor, A. M. (2020). The effects of quasi-random monetary experiments. *Journal of Monetary Economics*, 112:22–40.

Jordà, Ò., Singh, S. R., and Taylor, A. M. (2024). The long-run effects of monetary policy. *Review of Economics and Statistics*, pages 1–49.

Jordà, Ò. and Taylor, A. M. (2016). The time for austerity: estimating the average treatment effect of fiscal policy. *The Economic Journal*, 126(590):219–255.

Jordà, Ò. and Taylor, A. M. (2024). Local projections. Technical report, National Bureau of Economic Research.

Känzig, D. R., Konradt, M., Wang, L., and Zhang, D. (2025). Green business cycles. Technical report, National Bureau of Economic Research.

Känzig, D. R. and Williamson, C. T. (2025). Unraveling the drivers of energy-saving technical change. Technical report, National Bureau of Economic Research.

Keuschnigg, C. and Stalenis, G. K. (2025). Energy saving innovation, vintage capital, and the green transition. Technical report, CESifo Working Paper.

León-Ledesma, M. A. and Satchi, M. (2019). Appropriate technology and balanced growth. *The Review of Economic Studies*, 86(2):807–835.

Lewis, R. and Reinsel, G. C. (1985). Prediction of multivariate time series by autoregressive model fitting. *Journal of multivariate analysis*, 16(3):393–411.

McDonald, K., Pinardon-Touati, N., and Walsh, C. (2026). Growth, firm scale, and the energy intensity of production. Working paper, Columbia University.

Myhre, G., Alterskjær, K., and Lowe, D. (2009). A fast method for updating global fossil fuel carbon dioxide emissions. *Environmental Research Letters*, 4(3):034012.

Newell, R., Jaffe, A., and Stavins, R. (1999). The induced innovation hypothesis and energy-saving technological change. *Quarterly Journal of Economics*, 114(3):941–975.

Obstfeld, M. and Taylor, A. M. (2004). *Global capital markets: integration, crisis, and growth*. Cambridge university press.

Olea, J. L. M., Plagborg-Møller, M., Qian, E., and Wolf, C. K. (2024). Double robustness of local projections and some unpleasant varithmetic. Technical report, National Bureau of Economic Research.

Piger, J. and Stockwell, T. (2023). Differences from differencing: Should local projections with observed shocks be estimated in levels or differences? *Available at SSRN 4530799*.

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

Quinn, D., Schindler, M., and Toyoda, A. M. (2011). Assessing measures of financial openness and integration. *IMF Economic Review*, 59(3):488–522.

ONLINE APPENDIX

A1. Emissions intensity by country

Figure A.1: *Emissions intensity measured over real GDP (1 of 4)*

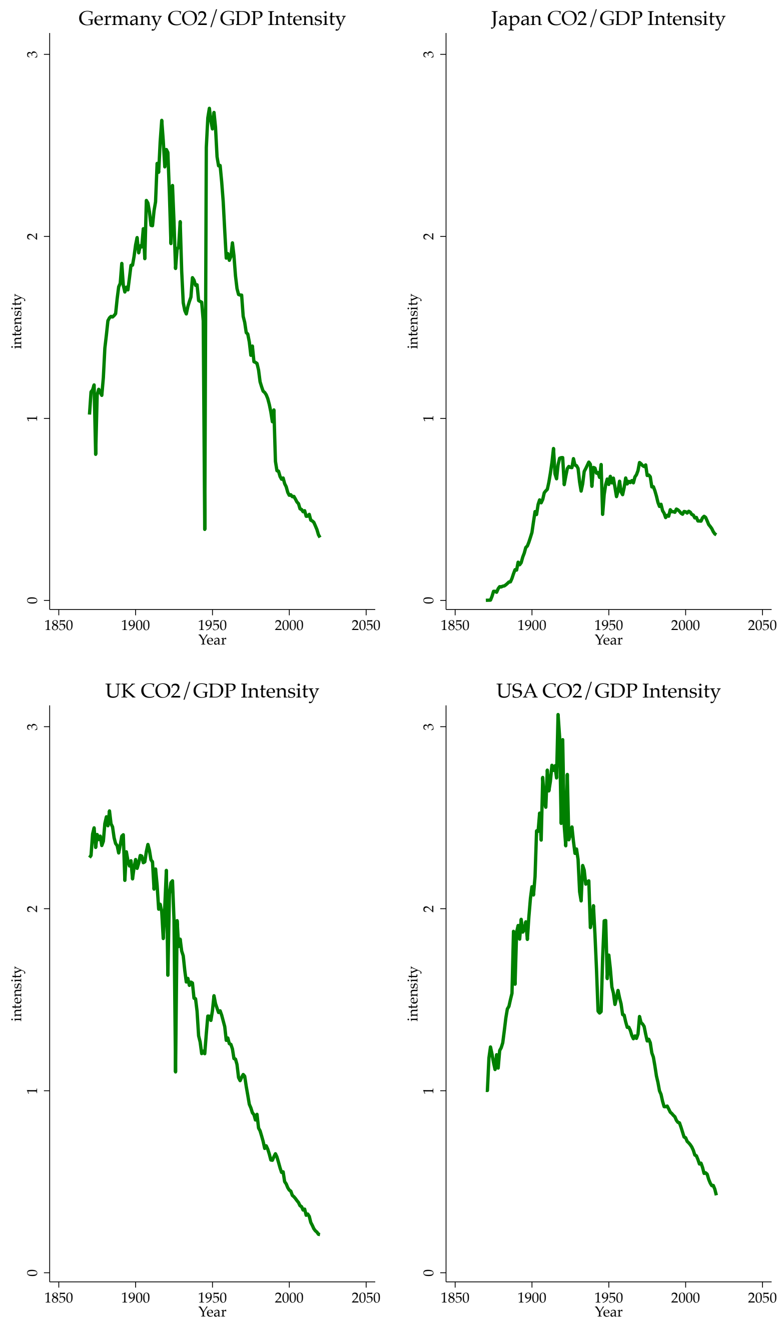

Notes: Total CO2 emissions for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. See text.

Figure A.2: *Emissions intensity measured over real GDP (2 of 4)*

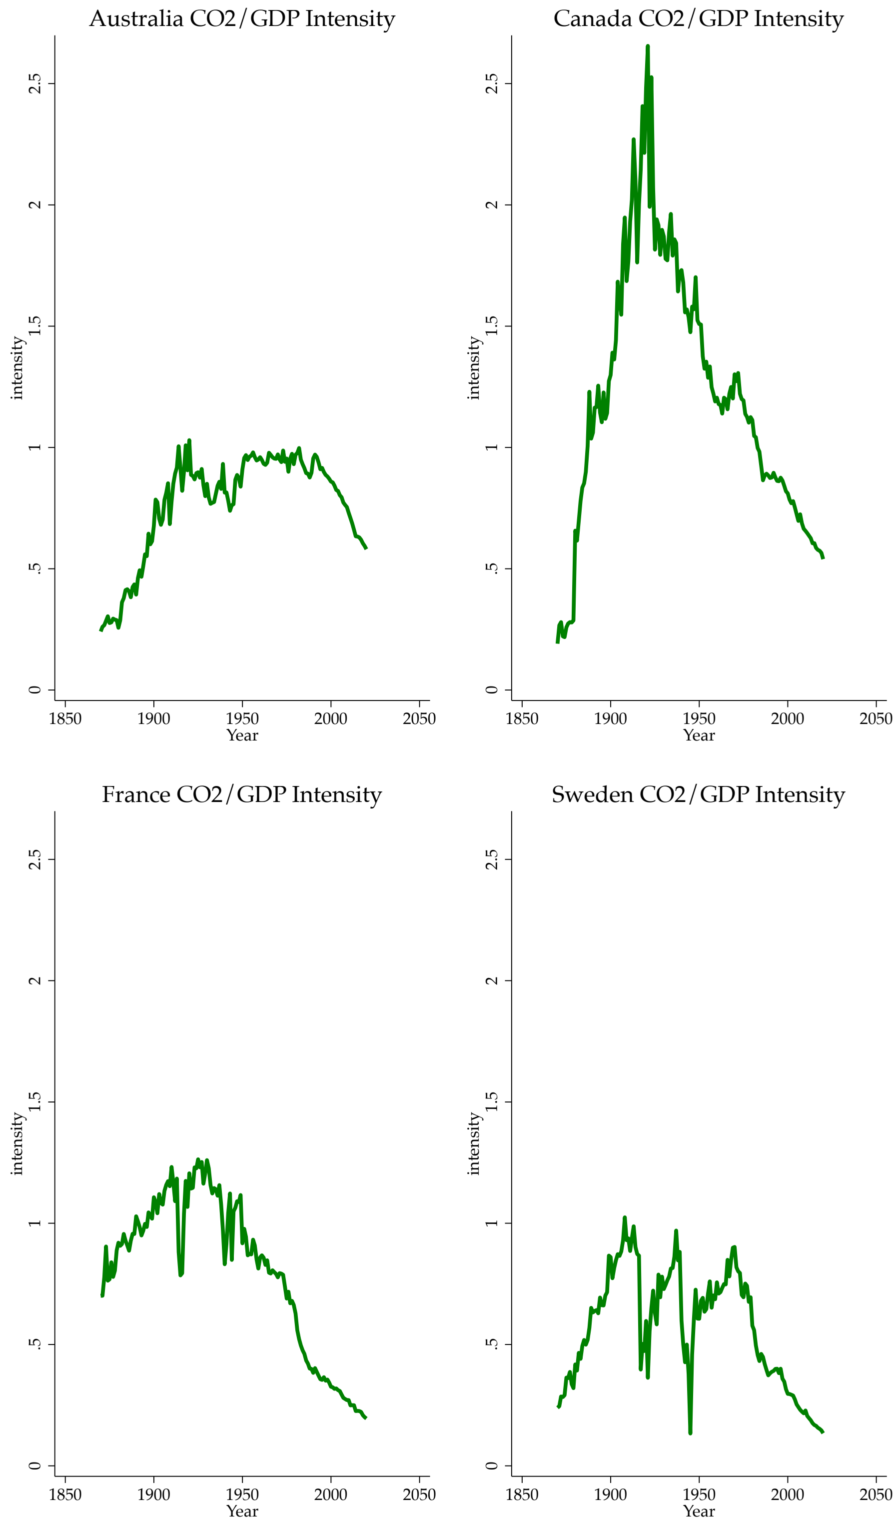

Notes: Total CO2 emissions for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. See text.

Figure A.3: *Emissions intensity measured over real GDP (3 of 4)*

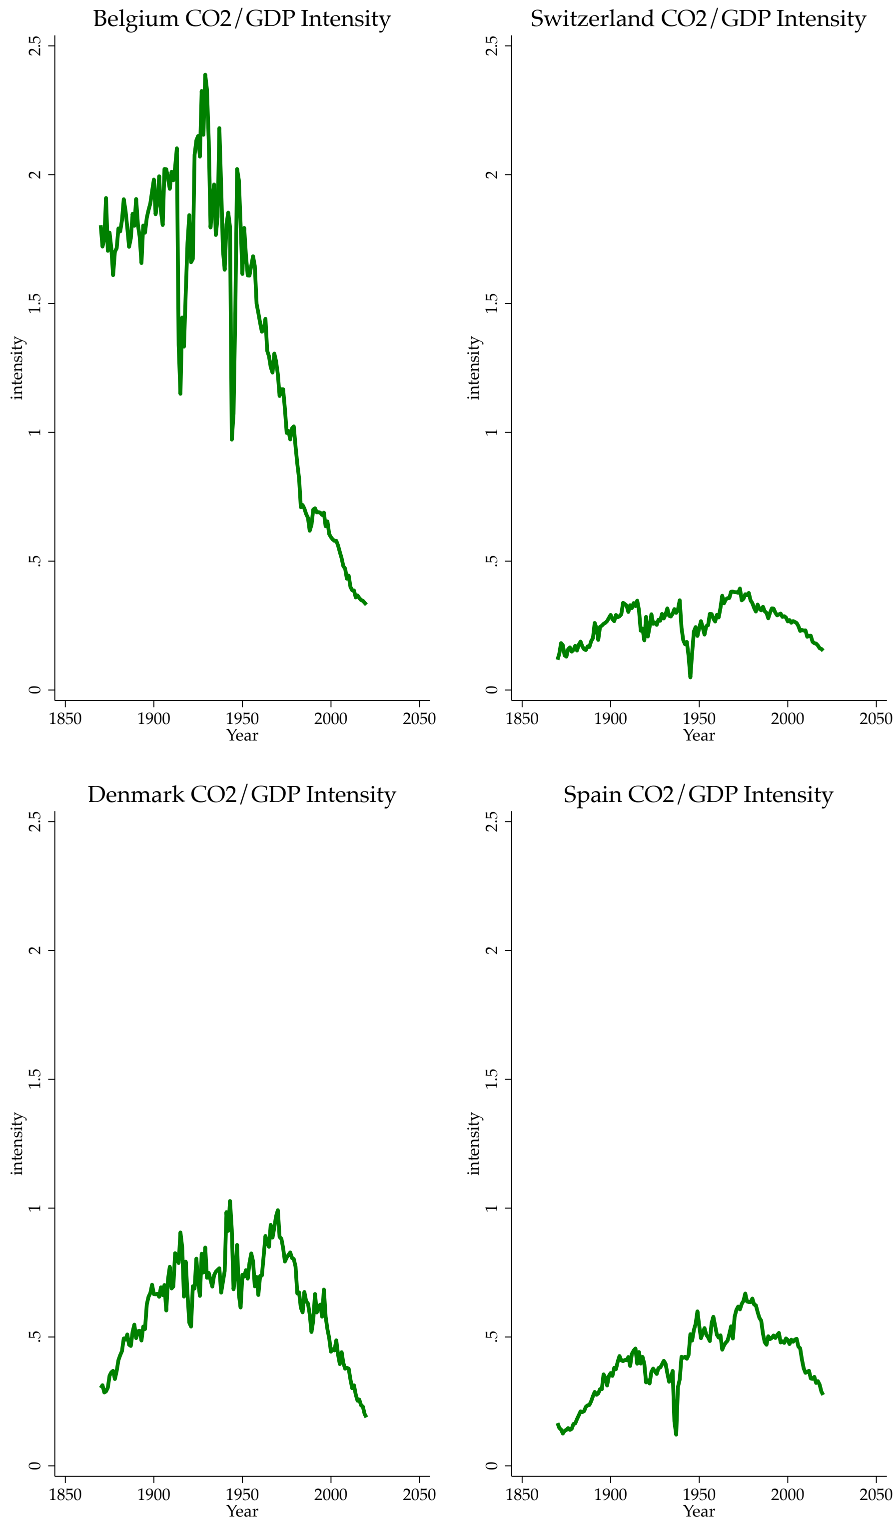

Notes: Total CO2 emissions for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. See text.

Figure A.4: *Emissions intensity measured over real GDP (4 of 4)*

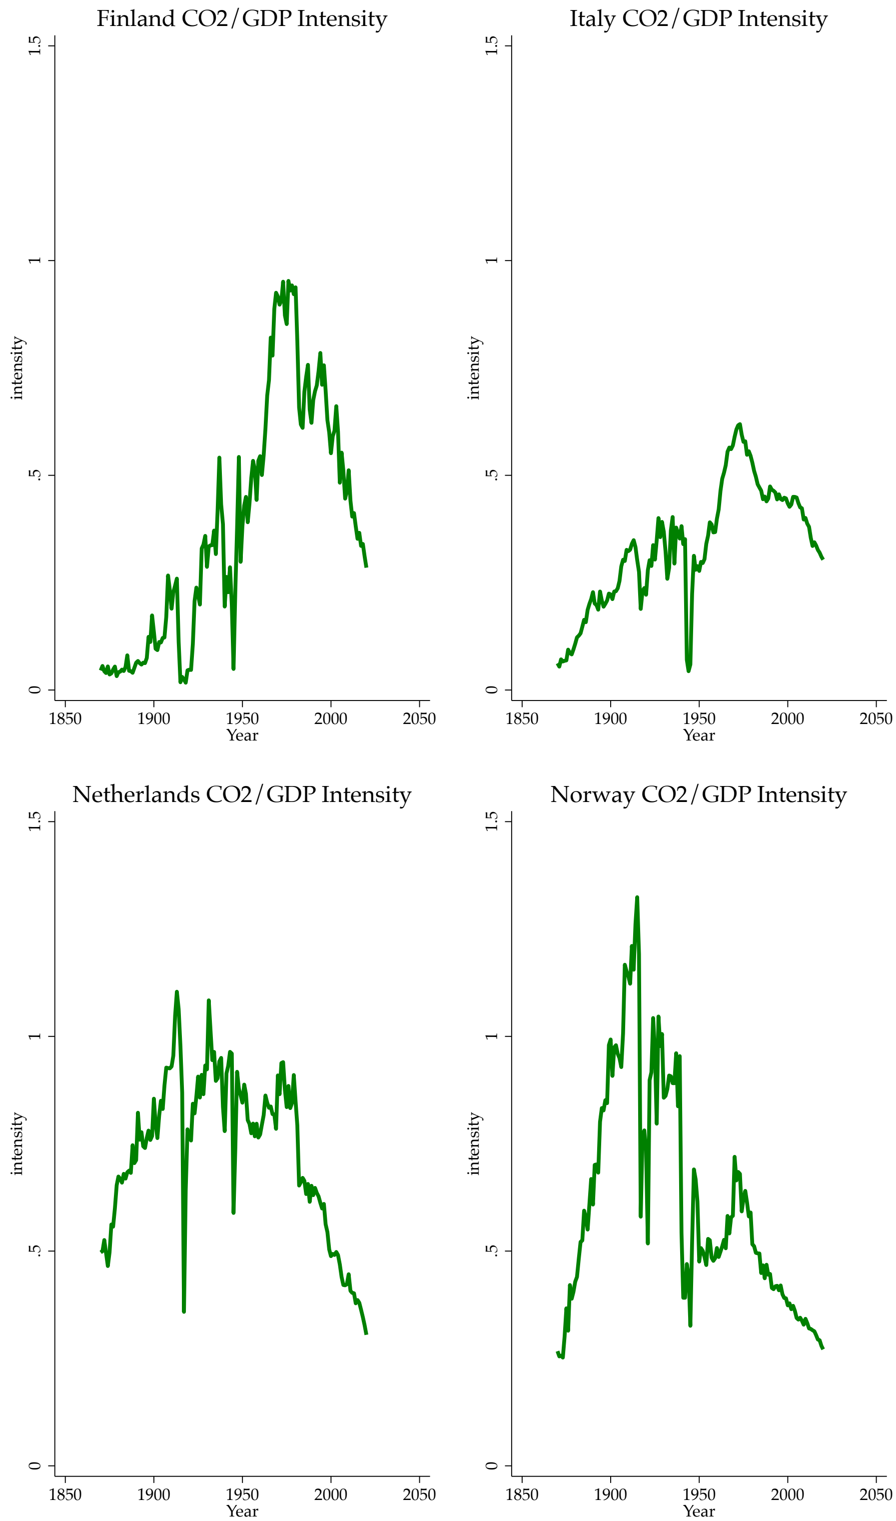

Notes: Total CO2 emissions for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. See text.

A2. Emissions and energy intensity by country

Figure A.5: *Emissions and energy intensity measured over real GDP (1 of 4)*

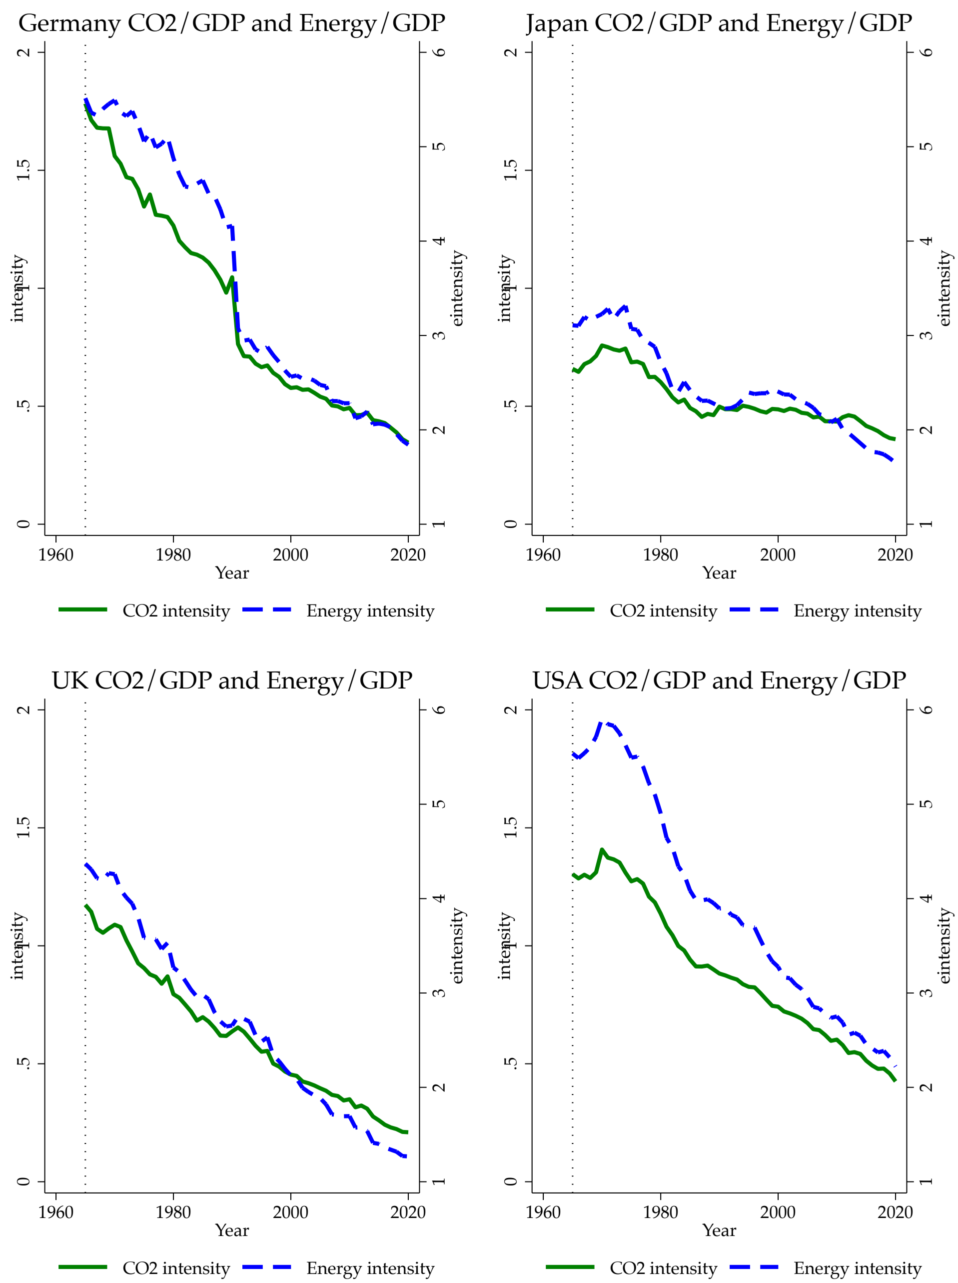

Notes: Total CO2 emissions and primary energy consumption for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. See text.

Figure A.6: *Emissions and energy intensity measured over real GDP (2 of 4)*

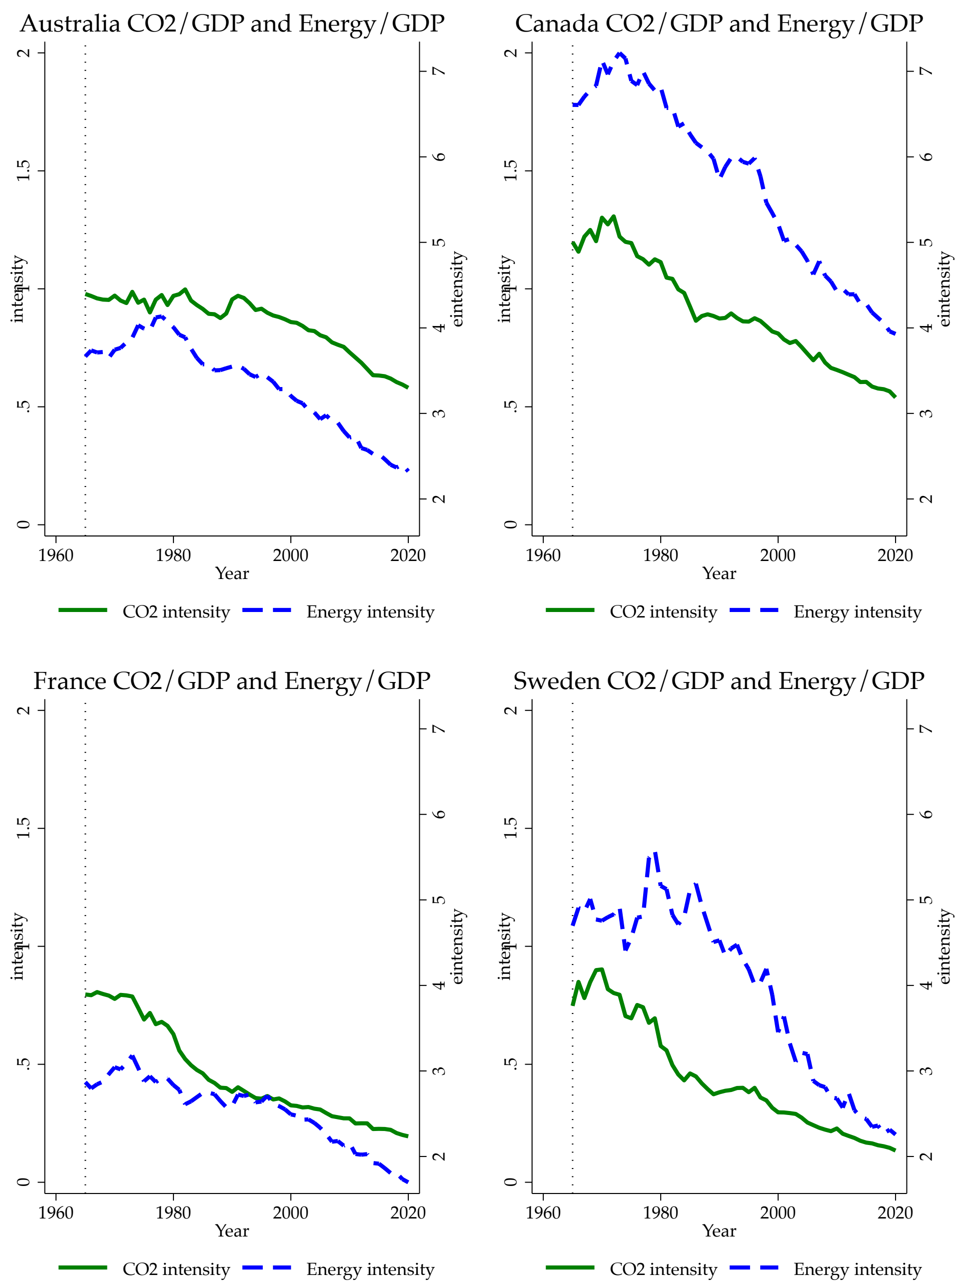

Notes: Total CO2 emissions and primary energy consumption for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. See text.

Figure A.7: *Emissions and energy intensity measured over real GDP (3 of 4)*

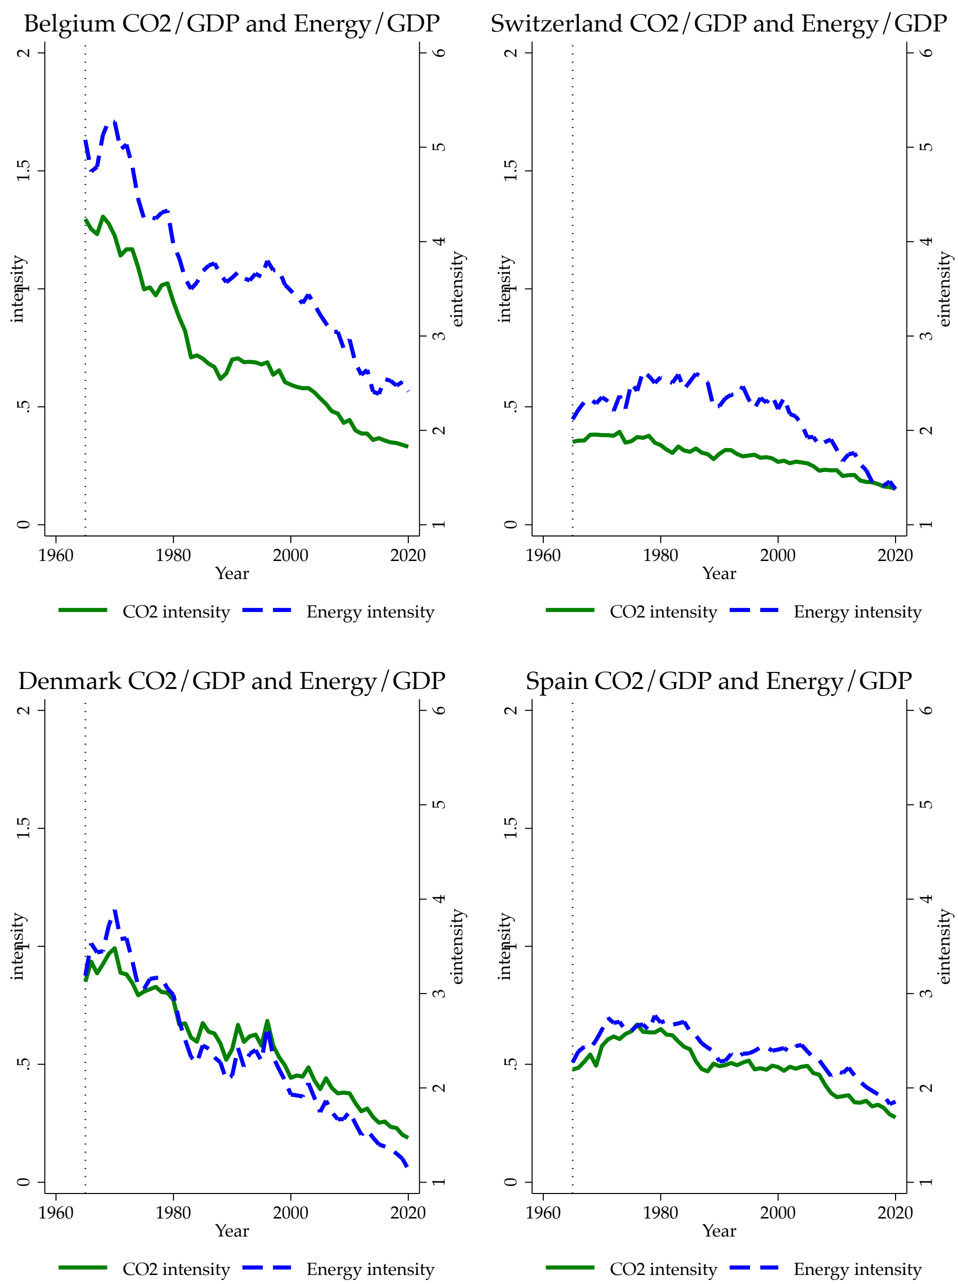

Notes: Total CO2 emissions and primary energy consumption for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. See text.

Figure A.8: *Emissions and energy intensity measured over real GDP (4 of 4)*

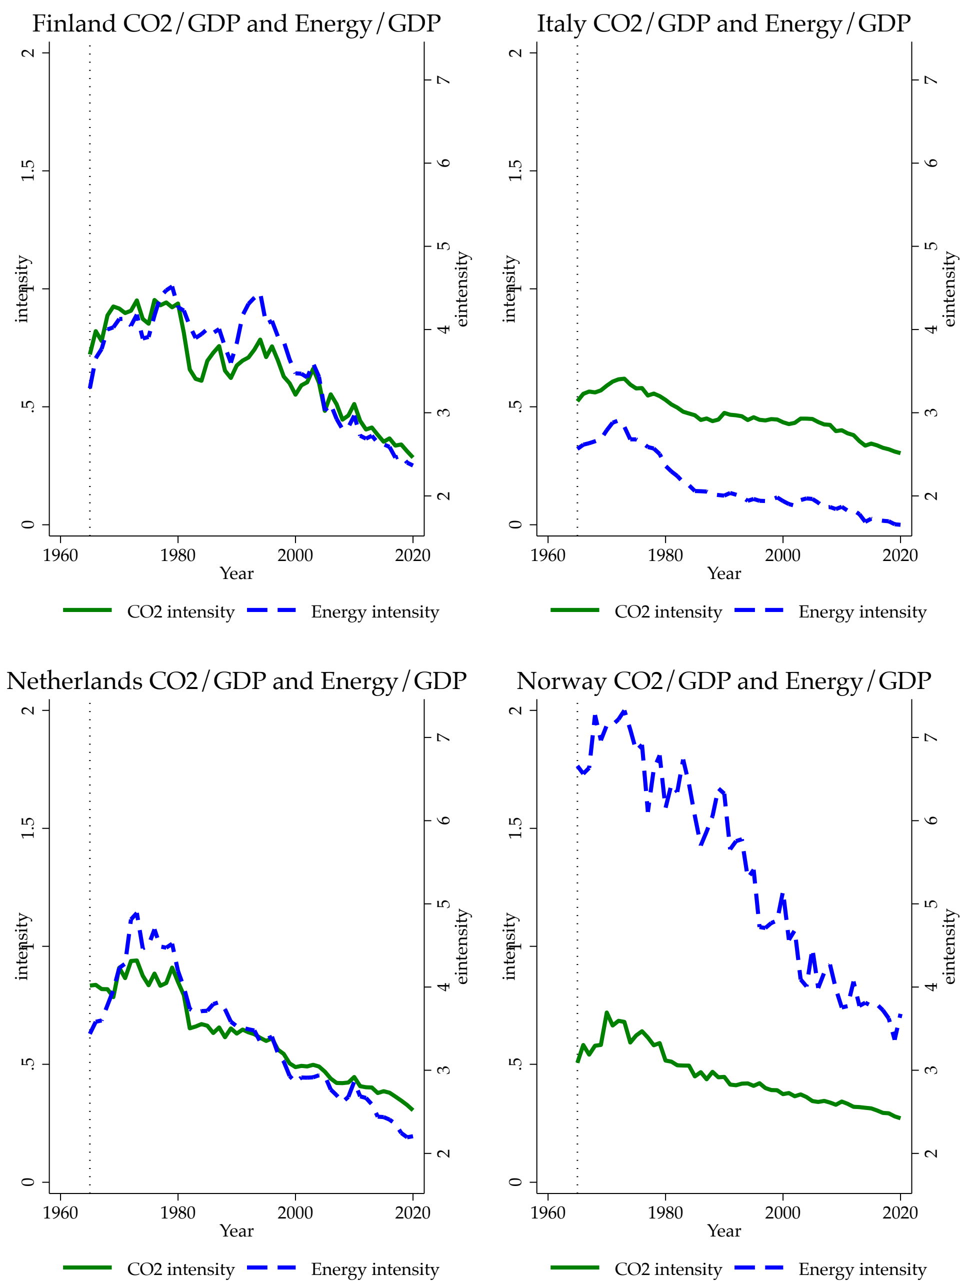

Notes: Total CO2 emissions and primary energy consumption for the 17 countries from the JST database divided by total real GDP in PPP adjusted 1990 dollars. See text.

A. Data Sources and Measurement of CO2 Emissions

A.1. Overview of CO2 emission data

This paper uses historical greenhouse gas and carbon emissions data compiled by Our World in Data (OWID), which provides country-level emissions estimates dating back to approximately 1750. The dataset captures CO2 emissions from various sources, including fossil fuel combustion (disaggregated by solid, liquid, and gas fuels), gas flaring, cement production, and international bunkers. This appendix details the data collection methodology, processing chain, and potential limitations.

A.2. Data Processing Chain

The carbon emissions data used in this project follows a multi-step processing chain that combines information from various institutional sources:

$$\text{Country reporting} \rightarrow \text{United Nations} \rightarrow \text{CDIAC-FF/BP} \rightarrow \text{Global Carbon Project} \rightarrow \text{Our World in Data}$$

Each step in this chain involves specific methodological choices and data transformations:

A.2.1 Primary Data Collection (Country Level)

Countries report emissions data to the United Nations following the Intergovernmental Panel on Climate Change (IPCC) guidelines (IPCC, 2006b,a), which outline two primary accounting approaches:

1. **Reference Approach**: A top-down method based on "apparent consumption" of fossil fuels, calculated using the following formula:

   $$\text{Apparent Consumption} = \text{Production} + \text{Imports} - \text{Exports} - \text{International Bunkers} - \text{Stock Change}$$

   This approach provides a relatively simplified estimate of emissions based on fuel supply rather than actual combustion. According to the IPCC guidelines, the Reference Approach is a straightforward method that uses a country's energy supply data to calculate CO2 emissions from fossil fuel combustion (IPCC, 2006a). This data is broken down by fuel source (i.e., coal, natural gas, etc.). It serves as a verification tool for the Sectoral Approach and functions as a quality assurance measure by comparing the two independent emission estimates.

2. **Sectoral Approach**: A bottom-up method based on measured fuel combustion, typically derived from national energy statistics. The IPCC guidelines define three tiers of increasing complexity:

- **Tier 1**: Uses default emission factors and basic activity data
- **Tier 2**: Uses country-specific emission factors
- **Tier 3**: Uses detailed facility-level data and/or sophisticated models

Most countries report emissions using both approaches, though the values typically show minor differences (generally less than 5%) due to methodological variations, statistical differences, inclusion of feedstocks, and differences in carbon stored in non-energy products (IPCC, 2006a).

A.2.2 United Nations Data Compilation

The United Nations compiles and standardizes the country-reported data, serving as the primary international repository for emissions information. The UN data follows standardized formats and undergoes quality control procedures before being made available for further analysis.

A.2.3 CDIAC-FF and BP Data Integration

The Carbon Dioxide Information Analysis Center - Fossil Fuel (CDIAC-FF) dataset, historically maintained by Oak Ridge National Laboratory and now by Appalachian State University (Gilfillan and Marland, 2021), primarily uses the UN reference approach data. This choice is largely driven by the greater historical availability of reference approach data, particularly for periods before comprehensive sectoral measurement systems were established.

According to Andrew and Peters, the history of CDIAC's emissions estimates goes back to 1973 with the work of Keeling (1973) and Rotty (1973), with CDIAC consistently producing national and global emissions estimates since 1999 (Andrew and Peters, 2021). CDIAC applies standard factors to apparent consumption of energy derived from UN energy data and includes emissions from flared natural gas.

For more recent years where UN data lags (typically by two years), emissions estimates from BP's Statistical Review of World Energy are incorporated to extend the dataset's temporal coverage (Myhre et al., 2009; BP, 2021). BP's data is primarily based on energy consumption statistics, which are then converted to emissions using standard conversion factors. Since BP only provides detailed data for major countries (with remaining countries grouped into regional categories), this approach introduces some uncertainty when applying regional growth rates to individual countries within those regions.

A.2.4 Global Carbon Project Processing

The Global Carbon Project (GCP) applies several important modifications to the CDIAC-FF data to create what they describe as "the best possible estimate of fossil CO2 emissions globally" (Andrew and Peters, 2021). Their philosophy prioritizes accuracy over methodological consistency, with the condition that both double counting and undercounting of emissions are avoided.

Major refinements made by the GCP include:

1. Replacement of CDIAC estimates with official estimates from developed countries (particularly Annex-1 parties to the UNFCCC) where these are deemed more accurate due to:
   - Use of significantly more detailed data and information
   - Expertise developed over many years
   - External auditing processes
   - Use of country-specific emission factors
2. Correction of CDIAC estimates that show clear disagreement with other reliable sources, including the International Energy Agency (IEA), which also uses a reference approach but with more detailed energy data and more extensive cross-checks
3. Use of more accurate final-year data where available, rather than extrapolations based on BP growth rates
4. Correction of implausible values in CDIAC data, such as negative emissions or unexplainable spikes
5. Country-specific corrections and improvements, including:
   - Comprehensive revisions for countries with significant oil and gas production/exports (e.g., Norway)
   - Addition of emissions from lime production in China (approximately 170 Mt CO2 in recent years)
   - Improved estimates for countries where the reference approach produces poor results (e.g., Indonesia)
   - Use of monthly data for countries like India to correct for fiscal year reporting issues
6. Complete replacement of CDIAC's cement process emissions with a more accurate methodology developed by Andrew (Andrew, 2019)
7. Special adjustments for international transport emissions in exceptional periods (such as during the COVID-19 pandemic)

The GCP also maintains continuous country definitions despite changing political boundaries, using various allocation methods to create unbroken time series for currently existing countries.

A.2.5 Our World in Data Compilation

Our World in Data compiles the GCP data along with additional metadata and contextual information to create the publicly available dataset used in this paper. OWID maintains the data in accessible formats with regular updates as new information becomes available.

A.3. Historical Data Sources (Pre-1960)

Historical emissions data, particularly for periods before 1960, derives primarily from fossil fuel production records rather than direct emissions measurements. Significant portions of these early data come from Etemad et al.'s "World Energy Production 1800-1985", which compiled country-year production estimates for coal, oil, and other fossil fuels from historical records (Etemad et al., 1991).

These production values are converted to emissions estimates using the reference approach methodology, with adjustments for estimated international trade where such information is available. However, it should be noted that the accuracy of these historical estimates decreases for earlier periods due to limitations in the historical record and changes in territorial boundaries.

According to Andrew and Peters, the earliest energy data in this chain come from Pollard (1980) for the UK, starting from 1750 (Andrew and Peters, 2021). The GCP dataset actually begins in 1750 rather than 1751 as previously reported by CDIAC, correcting a minor error in the original data transcription by Etemad and Luciani.

A.4. Methodological Variations and Limitations

Several methodological considerations should be noted when interpreting the emissions data:

1. **Methodological Heterogeneity**: The dataset prioritizes accuracy over methodological consistency, resulting in different calculation approaches across countries and time periods. While this increases the reliability of individual data points, it may introduce challenges for strict cross-country or temporal comparisons.

2. **Sectoral Resolution**: The dataset provides emissions by fuel source (coal, natural gas, oil), type (solid, liquid, gas), and certain production processes (cement production, gas flaring), but does not disaggregate emissions by consumption sector or industry. This is largely due to the historical availability of data in the early years, when UN data provided little information about sectoral energy consumption.

3. **Territorial Attribution**: Emissions are attributed to the territory where fossil fuels are combusted, rather than where the resulting goods or services are consumed. This approach does not account for "embodied carbon" in international trade.

4. **Geographic Definition Issues**: There are inconsistencies in the territorial definitions of some countries across different data sources, particularly for countries with overseas territories or complex political arrangements.

5. **Historical Uncertainty**: Earlier historical data (particularly pre-1900) carries greater uncertainty due to limited primary sources and changes in territorial boundaries over time.

6. **Extrapolation Uncertainty**: For the most recent 1-2 years, emissions are often estimated using growth rates derived from BP's energy data, which introduces additional uncertainty, particularly for smaller countries grouped into regional categories.

7. **Updates and Revisions**: The data undergoes periodic revisions as methodologies improve and new primary sources become available, potentially affecting historical estimates. For example, significant revisions occurred in China's coal consumption data following its third Economic Census.

A.5. Data Access

The dataset used in this analysis is publicly available through Our World in Data's GitHub repository at https://github.com/owid/co2-data/. Detailed variable descriptions and measurement units are documented in the codebook (owid-co2-codebook.csv).

B. The effect of oil prices on efficiency

This section replicates main figures Figure 6 and Figure 8 by adding how the response would change if oil prices went up by two standard deviations (50% above the mean). This response is shown as the purple dashed line along with the same response as standard errors as the figures in the text. The estimates are based on the specification in Equation 20.

Figure A.9: *Cumulative response of macroeconomic aggregates to a stance shock. 1890-2019*

(a) *Real GDP*

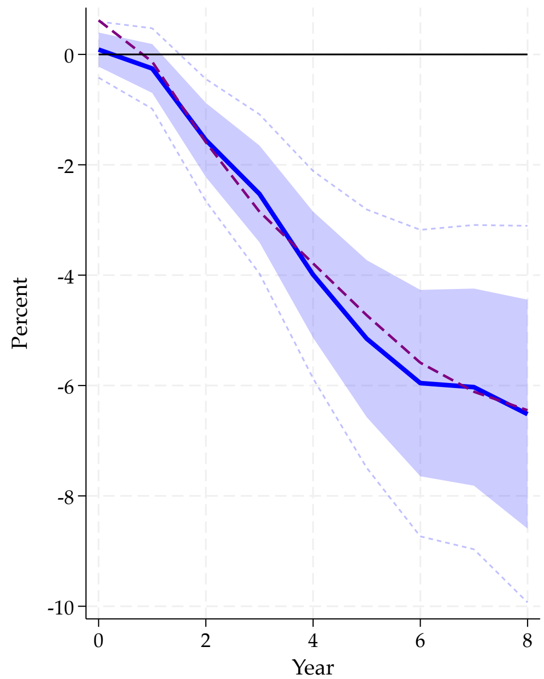

(b) *Real capital*

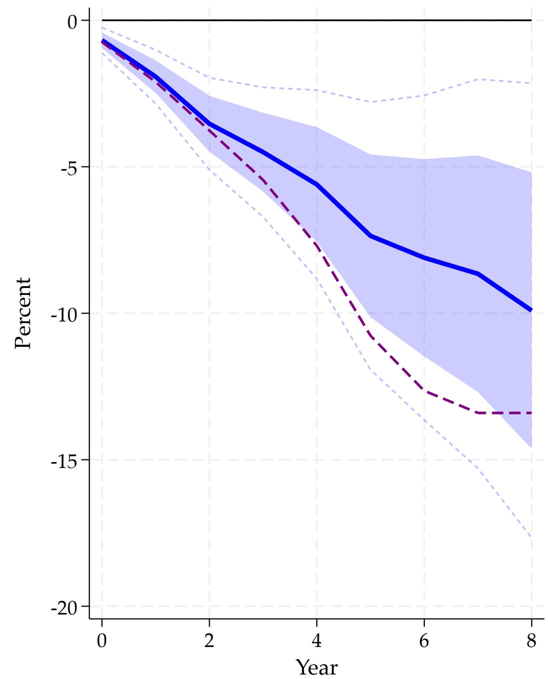

(c) *CO2 emissions*

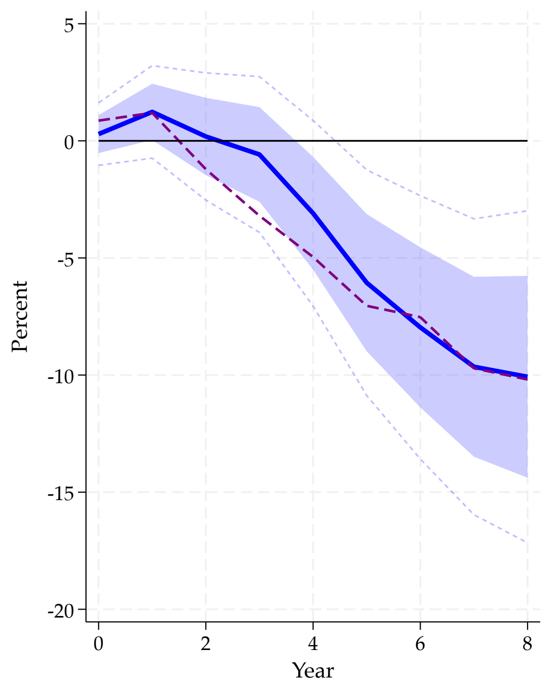

(d) *Short-term real interest rate*

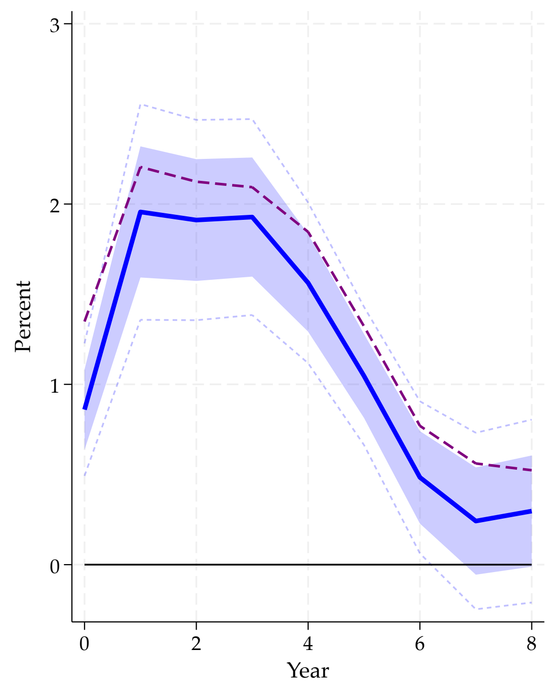

Notes: Cumulative responses to a shock in stance instrumented with the trilemma instrument. The response when oil prices are at the mean shown as solid blue line. The response when oil prices are 50% above the mean shown as dash-purple. 90% confidence bands and one standard deviation region shown for the response at the mean. Sample: 1890-2019 excluding World Wars. See text.

Figure A.10 digs deeper into the relationship between CO_2 emissions and the two types of capital available in our data: investment in buildings and structures versus investment in machinery. Our model does not make a distinction between machine and building capital so this breakdown in has no counterpart in the model.

Panel (a) displays the cumulative response of the ratio of CO_2 emissions relative to "buildings" capital whereas the panel (b) refers to the ratio of CO_2 emissions to "machine" capital. CO_2 emissions per unit of capital increase initially for both types of capital, though building capital, displayed in panel (a), seems to return back to its initial ratio after year 4, whereas the effect is somewhat more persistent for machine capital, displayed in panel (b). When oil prices are 50% above their mean, displayed in the figure with the purple dashed line, these effects are considerably more persistent.

Figure A.10: *Cumulative response of CO_2 emissions relative to capital. 1890-2019*

(a) *CO_2 emissions to building capital* **(b)** *CO_2 emissions to machine capital*

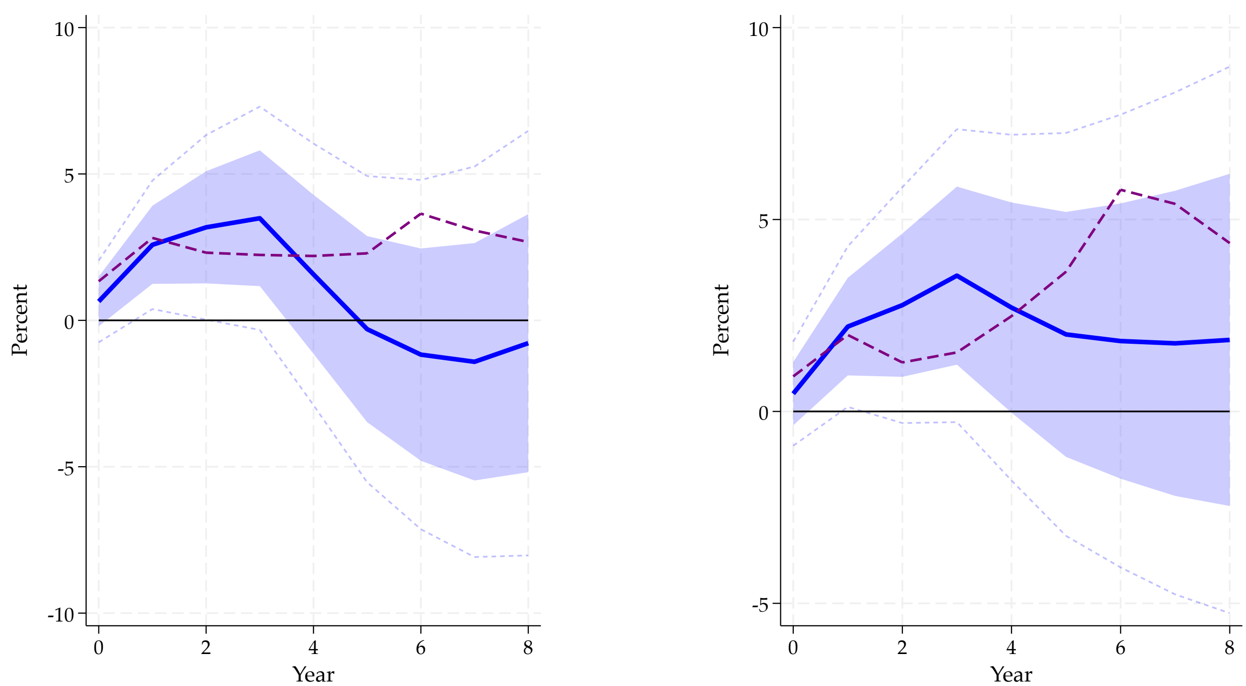

Notes: Cumulative responses to a shock in stance instrumented with the trilemma instrument. The response when oil prices are at the mean shown as solid blue line. The response when oil prices are 50% above the mean shown as dash-purple. 90% confidence bands and one standard deviation region shown for the response at the mean. Sample: 1890-2019 excluding World Wars. See text.

Figure A.11: *Cumulative response of macroeconomic aggregates to a stance shock. 1965-2019*

(a) *Real GDP*

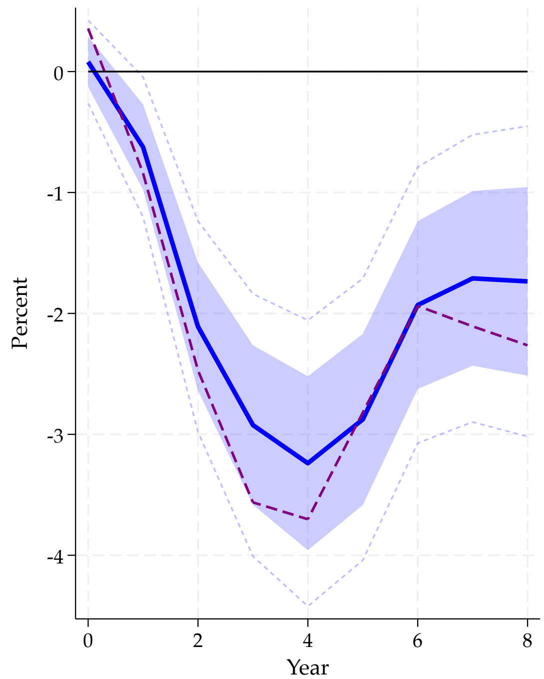

(b) *Real capital*

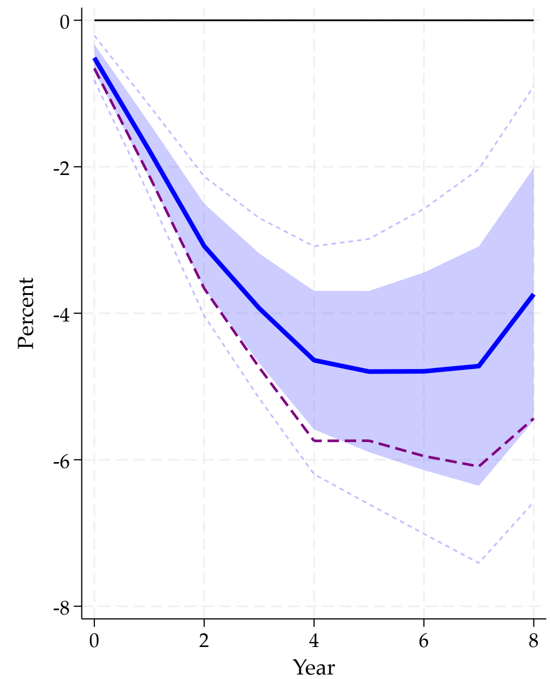

(c) *CO2 emissions*

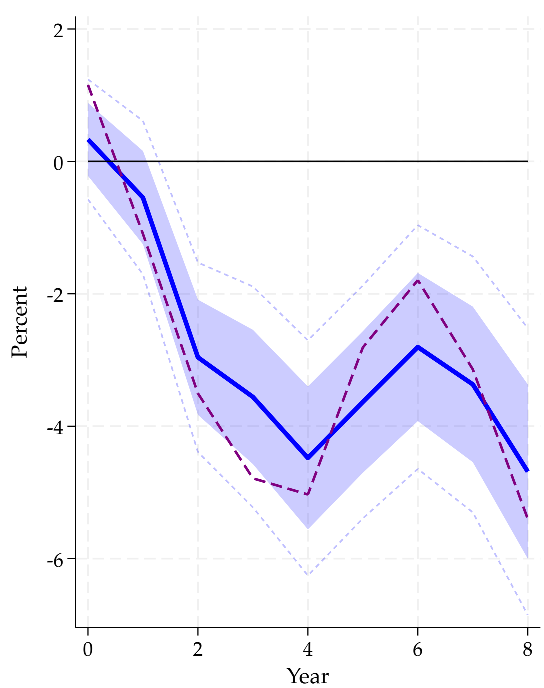

(d) *Short-term real interest rate*

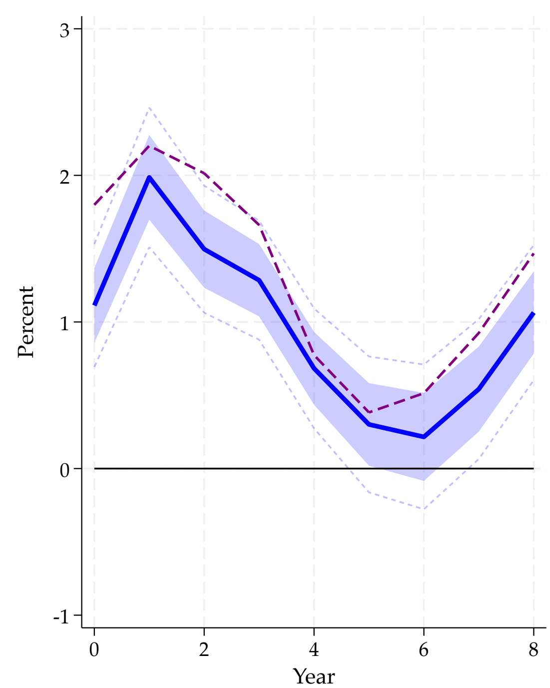

Notes: Cumulative responses to a shock in stance instrumented with the trilemma instrument. The response when oil prices are at the mean shown as solid blue line. The response when oil prices are 50% above the mean shown as dash-purple. 90% confidence bands and one standard deviation region shown for the response at the mean. Sample: 1965-2019 excluding World Wars. See text.

Figure A.12: *Cumulative response of CO2 emissions relative to capital. 1965-2019*

(a) *CO2 emissions to building capital* **(b)** *CO2 emissions to machine capital*

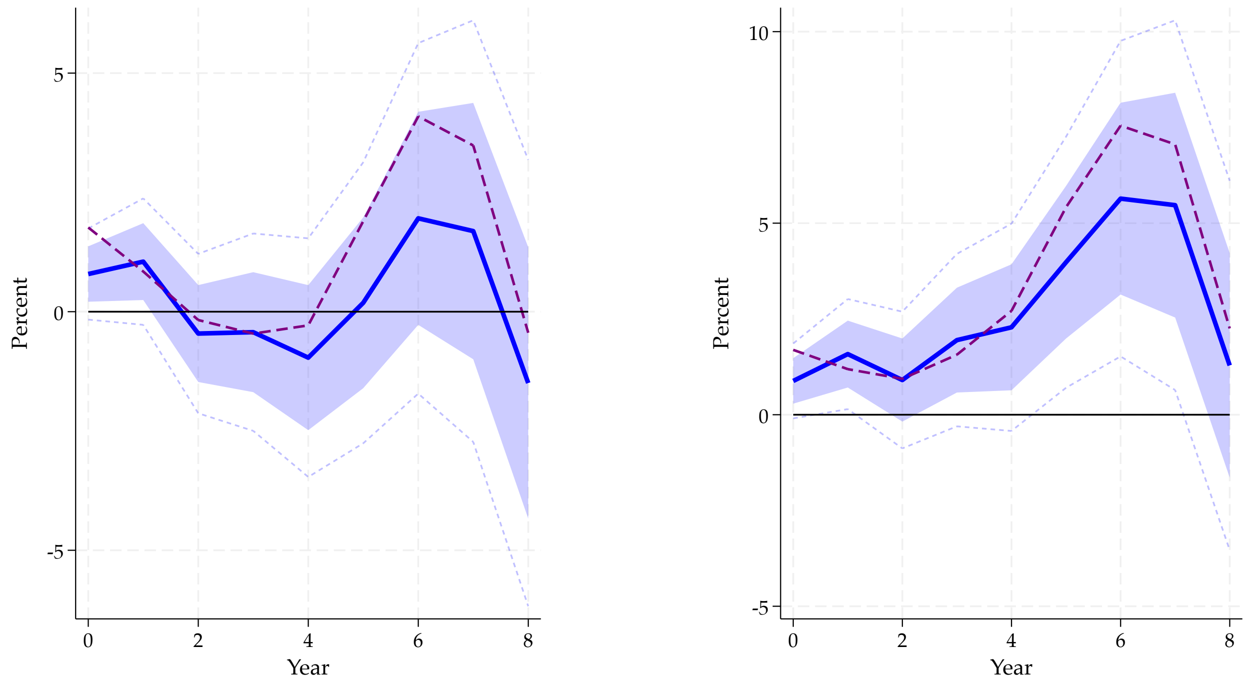

Notes: Cumulative responses to a shock in stance instrumented with the trilemma instrument. The response when oil prices are at the mean shown as solid blue line. The response when oil prices are 50% above the mean shown as dash-purple. 90% confidence bands and one standard deviation region shown for the response at the mean. Sample: 1965-2019 excluding World Wars. See text.

Figure A.13: *Energy intensity to capital*

(a) *Energy to building capital* **(b)** *Energy to machine capital*

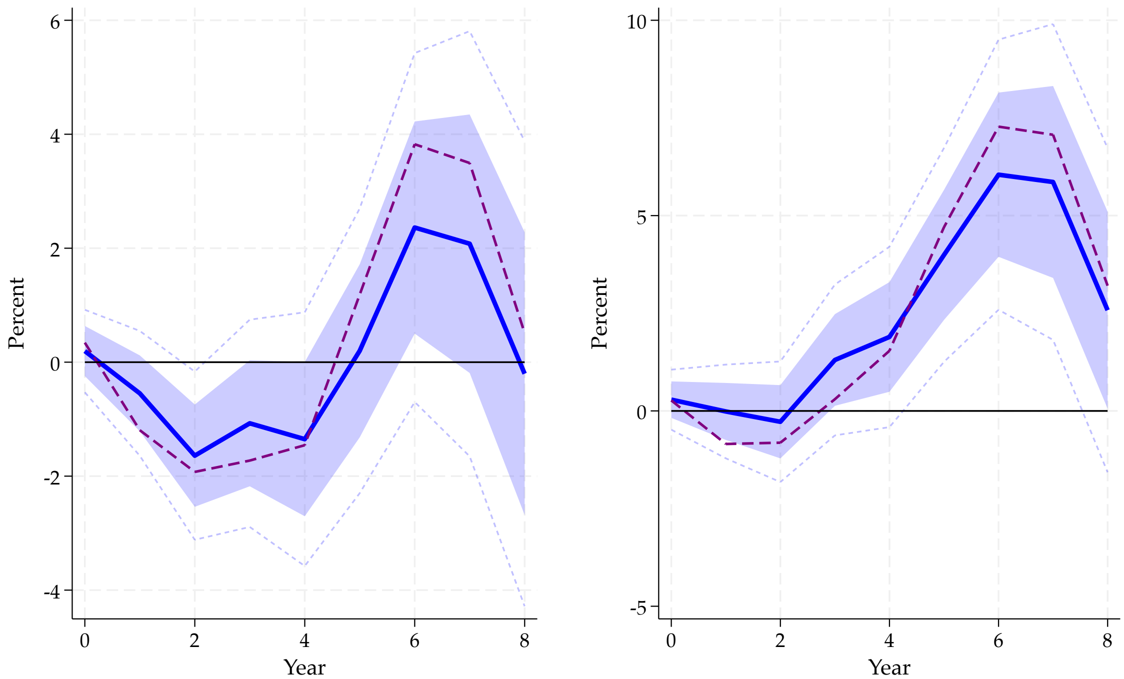

Notes: Cumulative responses to a shock in stance instrumented with the trilemma instrument. 90% confidence bands and one standard deviation region shown for the response at the mean. Sample: 1965-2019 excluding World Wars. See text.