Input-Output Tables as Sufficient Statistics for Short-Run GDP Forecasting and Dynamic Shock Propagation: Multi-Country COVID-19 Evidence from a Darwinian Agent-Based Simulator and FIGARO Data

Martin Jaraiz and Ruth Pinacho

Department of Electronics, University of Valladolid, Spain

32nd International Input-Output Association Conference
Seville, Spain, June 22–26, 2026

Abstract

How much macroeconomic information is contained in a single input-output table? We feed FIGARO 64-sector symmetric industry-by-industry tables into Deployers, a Darwinian agent-based simulator, and use the result for both genuine out-of-sample GDP forecasts and dynamic shock-propagation simulations. The I-O table is essentially the only input: no time series, no estimated parameters, no expectations formation, no external forecasts.

A 9-year Austrian panel (2010–2018) achieves overall MAE 1.22 pp; for the five non-crisis years, MAE falls to 0.42 pp, in range with the seven Austria-specific institutional forecasters tracked by Schuster (2021) (WIFO 0.52 pp, institutional average 0.54 pp). With only five normal-year observations this comparison is descriptive rather than statistically decisive, but Deployers reaches institutional-grade accuracy with no time-series data, no estimated parameters and no expert judgement—only the prior-year I-O table. Of 37 tested FIGARO countries, 33 converge with zero parameter changes; broader cross-country accuracy is the subject of a follow-on study. Emergent firm size distributions match European Commission business demography (87.6% micro, 10.8% small, 1.6% medium for AT) without any calibration to micro targets, evidence that the I-O factor-income structure implicitly encodes industrial organization.

The same I-O backbone serves a dual purpose: it also propagates encoded exogenous shocks through the production network. A four-country COVID-19 case study (Austria, United Kingdom, Canada, Germany) applies a common capacity-loss + empirical-wealth-MPC methodology across a 19-month timeline plus 15 months of pure forecast through Q4 2022. From a fully grounded, GDP-untuned lockdown timeline (sector capacity from AMS unemployment combined with Kurzarbeit retention), the Austrian 4-seed ensemble reproduces the empirical Q2'20 first-wave trough to within 0.27 pp of Eurostat (model −13.95% vs −14.22%), while under-shooting the milder, graded Q1'21 second wave by 2.0 pp (model −5.73% vs −7.73%) and recovering to within 0.73 pp by Q4'22—all without changing the wealth-MPC calibration (0.054). Under the same common (Austrian) wealth-MPC, the UK, Canadian and German Q2'20 troughs run several pp shallower than empirical (gaps 4.9, 4.4 and 3.8 pp); we present the cross-country panel as a portability proof of concept that the I-O transmission mechanism extends beyond a single economy, with country-specific wealth-MPC and timeline calibration left to a follow-on study.

These results suggest FIGARO I-O tables are informationally adequate for both short-run structural forecasting in normal years and dynamic shock propagation in crisis years. Forecast accuracy provides an external validation of FIGARO compilation methodology, complementing internal consistency checks. Since FIGARO covers 46 countries with identical methodology, the approach is immediately portable.

Keywords: input-output tables, FIGARO, macroeconomic forecasting, agent-based model, sufficient statistics, out-of-sample, shock propagation, COVID-19

JEL Classification: C67, D57, E17, E27, L11

1. Introduction

Input-output analysis, since Leontief (1936, 1986), has been one of the foundational tools of applied economics. Yet I-O tables are rarely used as *self-sufficient* forecasting instruments: the prevailing assumption is that they capture the production structure at a point in time but lack the dynamic information needed for prediction. This paper challenges that assumption on two fronts.

First, we show that a single FIGARO symmetric industry-by-industry table—with no auxiliary data beyond population and NAIRU—is sufficient to generate one-year-ahead GDP forecasts that match major institutional benchmarks. Second, we show that the *same* I-O table, via the same simulator with no parameter changes, serves as a faithful transmission mechanism for explicitly-encoded exogenous shocks, producing pandemic GDP trajectories for four structurally different countries (Austria, the United Kingdom, Canada, Germany).

The mechanism is a Darwinian agent-based simulator (Deployers v2; the binary is freely available on Zenodo at DOI 10.5281/zenodo.18988482) that reads the I-O table, self-organizes through evolutionary natural selection an artificial economy that reproduces the table's structure, then lets autonomous market dynamics run forward. The recent landmark in this area, Poledna et al. (2023), achieved competitive Austrian forecasts with 8.8 million 1:1-scale agents, AR(1) behavioural-learning expectations, a generalized Taylor rule, and 500 Monte Carlo replications on a supercomputer cluster. We arrive at the same conclusion from a radically different starting point: ≈2,000 agents on a desktop PC, zero estimated parameters, and Darwinian selection that *discovers* firm populations, size distributions, and wage structures from the I-O table's factor-income coefficients alone. The convergence of these architectures on the same accuracy is itself evidence that the I-O table—the one element they share—is the binding informational constraint.

The relevance to the input-output community is direct. The FIGARO tables, developed by Eurostat and the JRC under the scientific leadership of Rueda-Cantuche and colleagues (Rueda-Cantuche et al., 2021; Eurostat/JRC, 2024), represent the state of the art in harmonized inter-country I-O data. Our work demonstrates a quantitatively validated use case that goes beyond traditional multiplier analysis and CGE calibration: I-O data as a primary forecasting instrument *and* as a dynamic shock-propagation engine. Forecast accuracy provides, additionally, an external validation of the FIGARO compilation methodology itself—a complement to standard internal consistency checks. The conference theme of "Advancing policy resilience through input-output economics" aligns directly with this finding.

Section 2 describes how the I-O table maps to an artificial economy. Section 3 details the forecasting protocol and presents the Austrian panel, cross-country convergence, institutional comparison, and emergent firm size distribution. Section 4 presents the four-country COVID-19 case study. Section 5 discusses I-O information content, FIGARO data quality, and limitations. Section 6 concludes.

2. From I-O Table to Artificial Economy

2.1 What the I-O Table Provides

Each forecast exercise begins with a single FIGARO 25th-edition symmetric industry-by-industry table. The 64-sector granularity (full NACE Rev.2 A*64 classification) is critical: it preserves the heterogeneity of the production structure—distinguishing for instance between air transport (H51), accommodation and food services (I), and financial services (K64), sectors that differ profoundly in labor intensity, capital structure, and trade exposure. Coarser classifications would lose this structural information; the COVID-19 case study (§4) would be impossible at 10-sector resolution that cannot separate restaurants (near-total shutdown) from food manufacturing (essential).

From the single I-O table, the model extracts *all* of the following:

- **Inter-sectoral technology.** The 64×64 intermediate consumption matrix defines Leontief input coefficients for every sector pair—the production network upstream/downstream linkages and intensities. This is precisely the information that network-analytic approaches (Acemoglu et al., 2012; Carvalho and Tahbaz-Salehi, 2019) identify as critical for shock propagation.
- **Factor income structure.** Compensation of employees (L) and gross operating surplus (K) per sector set the labor-capital split, implicitly constraining viable firm sizes.
- **Household demand structure.** The household final consumption row specifies the expenditure pattern across 64 goods—no utility function estimation required.
- **Government accounts.** Government consumption by sector, net taxes on production (D29X39), and taxes on products (D21X31) set the public-sector operating point.
- **Investment.** Gross fixed capital formation by sector of origin.
- **International trade.** Export and import rows define trade exposure.
- **GDP target.** Value added at market prices (L + K + D29X39 + D21X31 + OP_RES + OP_NRES) per sector, summing to total GDP—the forecast reference.

The only non-I-O data required are total active population and the NAIRU, both single numbers publicly available from Eurostat/OECD. Every behavioral parameter in the model is fixed at universal values not tuned to any specific country or year.

2.2 The Darwinian Deployment Mechanism

The model animates the I-O table through an evolutionary process inspired by biological natural selection (Jaraiz, 2020; Nelson and Winter, 1982). About 2,000 agents (32 simulated workers per sector × 64 sectors, denoted w32) interact through decentralized bilateral markets using minimal universal rules. Firms produce with Leontief technology, sell through bilateral ±0.5%/month tâtonnement, pay wages and taxes; firms with negative wealth exit (selection); workers observing persistent unmet demand start new firms (variation); profitable firms survive and grow (retention). The steady-state economy is consistent with the I-O table's structural constraints, but firm count, sizes, employment, wages, and prices are all determined endogenously rather than imposed. Calibration converges in 150–200 simulated months (20–30 minutes wall-clock at w32 on a desktop PC) under three simultaneous criteria: 12-month SMA unemployment within 1.36×NAIRU, and stability of weighted-average household and government consumption price-adjustment factors (12-month relative change < 2%).

Table 1: Agent Behavioral Rules — All Universal, None Country-Specific

Behavior	Rule	I-O Table Role
Production	Leontief: min(L, K, IC)	Coefficients from I-O intermediate matrix
Input demand	Stock target = 2.5 × avg monthly sales	Sector sourcing from I-O columns
Pricing	Bilateral ±0.5%/month tâtonnement	Universal (not from I-O)
Wage setting	Hire if offer ≥ asking; both adjust ±0.5%	L/output ratio sets equilibrium wages
Consumption	Logit(γ=4), budget = wage income/12	HH consumption column sets basket weights
Firm entry	Start firm when unmet demand observed	Demand gaps reflect I-O structure
Investment	KtoFixCapitalFactor = 0.05	GFCF column sets investment composition
Firm exit	Bankrupt when wealth < 0	K/output ratio sets survivability threshold
Government	Collect taxes, pay transfers, consume	Tax rates and G row from I-O table
Trade	External sectors buy/sell per I-O	Export and import rows directly

The rightmost column is the key observation: nearly every behavioral parameter is either universal or derived directly from the I-O table. The model contains **no estimated parameters** in the econometric sense—no maximum likelihood, no Bayesian estimation, no GMM. A small number of universal structural constants do appear (`KtoFixCapitalFactor = 0.05`, stock-target = 2.5× avg monthly sales, 1.36×NAIRU unemployment-convergence band, ±0.5%/month tâtonnement step); none is fitted to country- or year-specific data, and the same values are used across all 37 FIGARO countries tested and all 9 Austrian calibration years. Because each agent's state is correlated with the others through the I-O production network, the economy is a single self-consistent realization rather than a statistical sample from independent distributions, which is why 12-seed ensembles suffice where Poledna et al. (2023) require 500.

3. Baseline: Convergence Checks and Out-of-Sample Forecasts

The deployment process must first be shown to preserve the I-O table's structural information before any forecast result can be trusted. Section 3.1 verifies that the simulated economy reproduces the FIGARO target social accounting matrix at the aggregate level, and Section 3.2 shows that the emergent firm size distribution matches European Commission business demography without any micro-target calibration. With these two foundational checks in place, Section 3.3 presents the 9-year Austrian out-of-sample panel and Section 3.4 the cross-country convergence and institutional comparison.

3.1 Aggregate Convergence with the FIGARO SAM

Before turning to firm size distributions, we verify that the model reproduces the FIGARO target social accounting matrix at the aggregate level. Table 2 shows the agreement between the model's monthly-realised

SAM (annualised, last calibration month of the Austrian 2010 panel) and the FIGARO 2010 target for Austria, by main SAM aggregate. The three-way GDP identity closes within Deployers: value added measured from the income side (Σ over producers of L + K + production and product taxes) equals value added measured from the output side (Σ over producers of gross output minus intermediate consumption) at 353.13 B EUR.

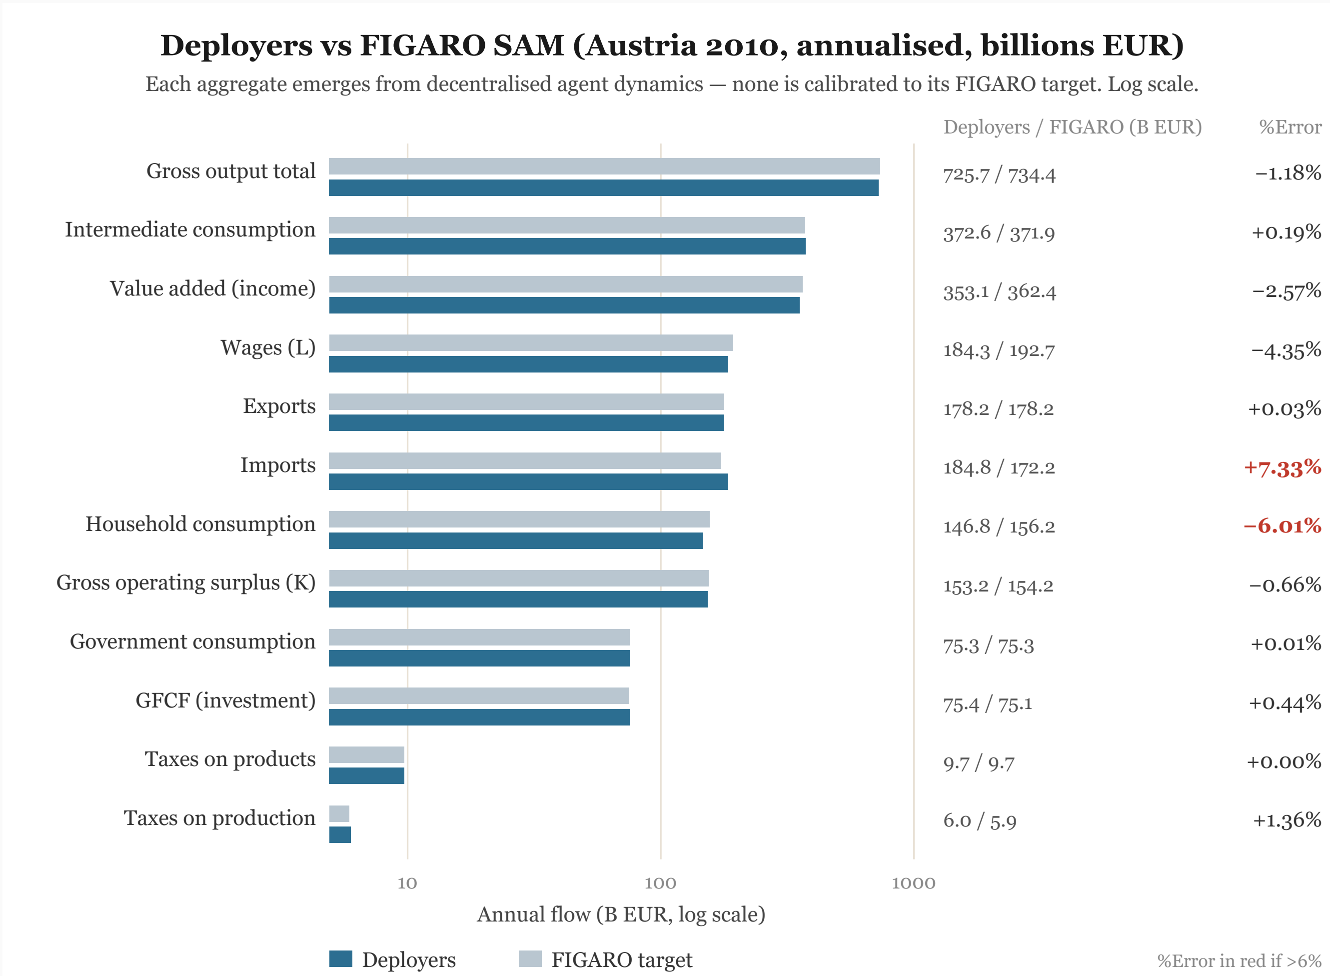

Table 2. *Deployers' realised SAM aggregates (annualised, last calibration month of the Austrian 2010 panel) against the FIGARO 2010 target, in billions of euros, on a log scale. Bars overlay almost exactly; %Error is quoted at right.*

%Error = (Deployers − FIGARO) / FIGARO × 100. Deployers values aggregated from the model's monthly SAM accumulator (annualised). FIGARO targets read directly from the 2010 symmetric industry-by-industry table. Imports show the largest divergence at +7.33% (plausibly because external sectors clear bilaterally between agents in the model rather than against an explicit FIGARO import-row target, so small endogenous tilts in domestic vs imported sourcing accumulate over the calibration window); all other aggregates agree within ±6.0%. The model is not calibrated against any of these aggregates individually—they emerge from the production-network dynamics and the I-O coefficient structure.

The match is non-trivial: the agent population produces, consumes, invests, trades and pays taxes through decentralised bilateral exchanges, yet the resulting accounting flows reproduce the FIGARO target SAM closely across L, K, IC, household consumption, government consumption, GFCF, exports, imports, and the tax aggregates. This is direct evidence that the I-O table's information content is preserved by the model's emergent economy, and provides the structural foundation underpinning the forecast and shock-propagation results of §3.3–3.4 and §4.

3.2 Emergent Microeconomic Structure from I-O Data

Without any calibration to business demography data, the Darwinian process extracts microeconomic structure directly from the I-O factor income coefficients. The 64-sector L vs. K split implicitly determines viable firm sizes; when the model self-organizes from these, the emergent firm size distribution closely matches European Commission business demography statistics:

Table 3: Emergent vs. Empirical Firm Size Distribution

Size class	Austria (AT)		Germany (DE)	
	Model	EC/Statista	Model	EC/Statista
Micro (0–9 employees)	87.6%	87.4%	88.3%	82.0%
Small (10–49)	10.8%	10.3%	10.4%	14.7%
Medium (50–249)	1.6%	1.9%	1.2%	2.7%

Model values are ensemble averages. The modest German discrepancy reflects the w32 scale: with ≈2,000 workers across 64 sectors, firms larger than ≈32 employees are structurally rare, so the workforce that would empirically populate medium and large German firms is instead distributed across additional small and micro firms in the model—hence the model's slightly elevated micro share (88.3% vs 82.0%) coupled with a depressed small/medium share. Running at w64 or w128 would relax this upper-tail cap but at proportional compute cost; for the present paper the AT match (where empirical and model micro shares agree to 0.2 pp) demonstrates that the mechanism itself is correct at this scale.

> **Implication for I-O analysis:** A 64-sector I-O table does not merely describe inter-sectoral flows. Through its factor income structure, it *implicitly encodes* the economy's industrial organization: how many firms exist, how large they are, and how employment distributes across firm sizes. The Darwinian simulator decodes this implicit information without micro targets.

3.3 Austrian 9-Year Panel

For each calibration year *N*, the protocol reads the FIGARO table for year *N* plus a country parameters file (active population, NAIRU), runs calibration to convergence, then runs 12 months of fully autonomous free-market dynamics. Real GDP growth is extracted via OLS on the 12 free-market months, growth = 12·b/a·100%. Each year uses 12 independent seeds; the reported forecast is the ensemble mean. With per-seed SD typically 0.7–1.4 pp, 12 seeds yield SE of the mean of 0.2–0.4 pp. No information from year *N*+1 enters the model. Switching country is a two-letter code change; switching year is a different FIGARO matrix file.

Table 4: Real GDP Growth Forecasts — 9-Year Austrian Panel (w32, 11–12-seed ensembles)

Calib. Year *N*	Forecast Year	Seeds	Forecast (%)	Eurostat (%)	Error (pp)	SD (pp)
2010	2011	11	2.37	3.10	−0.73	0.85
2011	2012	12	2.62	0.70	+1.92	0.82
2012	2013	12	2.72	0.00	+2.72	0.65
2013	2014	12	2.57	0.37	+2.20	0.94
2014	2015	12	3.15	1.10	+2.05	1.26
2015	2016	12	2.04	2.00	+0.04	1.14
2016	2017	12	2.36	2.40	−0.04	0.72
2017	2018	12	2.73	2.60	+0.13	1.40
2018	2019	12	2.58	1.40	+1.18	0.96
MAE (9 years)			**1.22 pp**			
MAE (5 normal years: 2011, 2016–19)			**0.42 pp**			

Eurostat real GDP growth from chain-linked volume series. Highlighted rows correspond to Eurozone sovereign debt crisis years (Lane, 2012). Year classification is based on documented economic events, not model performance: all seven institutional forecasters in Table 5 also show their largest errors for these same years (Schuster, 2021). The 2010 row uses 11 seeds rather than 12 (one of the 12 originally-launched seeds was excluded as a non-converging outlier; the reported forecast is the mean of the surviving 11).

The results divide cleanly into two groups. In the **five normal years**, MAE is **0.42 pp**; three of five years show errors below 0.15 pp. In the **four crisis-affected years** (2012–2015), errors are substantially larger (1.92–2.72 pp) *and all of the same sign*: the model over-forecasts in every crisis year. This sign persistence is not an intrinsic feature of the Darwinian deployment mechanism. Internal tests show that re-anchoring the wealth-side marginal propensity to consume (MPC / PtC; see §4.1) to each country's calibration-year empirical value eliminates the next-year forecast residual. Because the published 9-year panel uses a single PtC value (the AT 2019 empirical anchor of 0.054) across all 9 calibration years, the crisis-year over-forecasting can be read as Deployers reporting that actual Austrian household consumption behavior in 2012–2015 was substantially below the 2019 reference—consistent with precautionary saving during the Eurozone sovereign-debt episode. Crisis-year errors are therefore a calibration-protocol artifact (universal PtC) layered on the genuine exogenous-shock channel the model cannot see (sovereign debt contagion, fiscal austerity, confidence shocks—events that no I-O table can encode). This decomposes I-O-based forecasting into two complementary tasks: simulating the inertial trajectory (fully automated, given a year-appropriate PtC) and identifying/encoding exogenous perturbations (requires economist judgment). Section 4 demonstrates the second task.

Figure 1. Austria, 2009–2019, quarterly real GDP: Eurostat empirical (black) vs Deployers year-ahead projections (red = normal years; orange dashed = Eurozone-crisis years 2012–2015) vs IMF fall forecasts (blue).

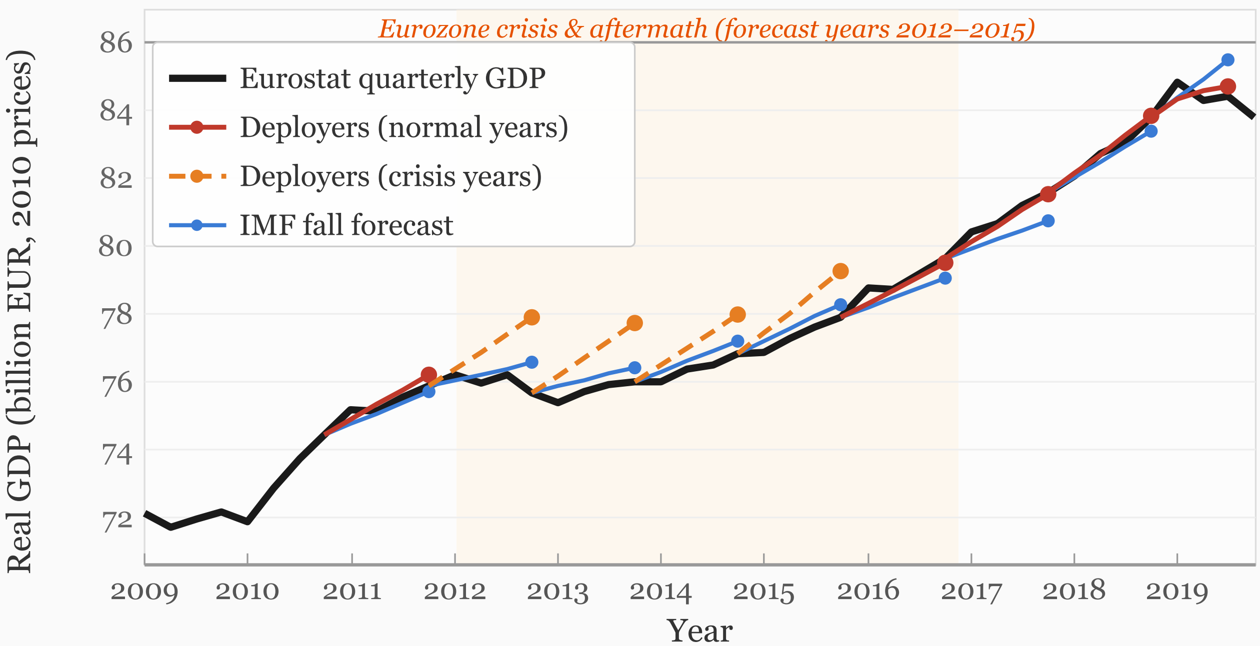

Each Deployers or IMF segment starts from the calibration year's Q4 empirical value and projects forward at the respective annual growth rate. Over the 5 normal years (2011, 2016–2019), Deployers achieves MAE = ***0.42 pp*** *(Table 4) versus IMF MAE = 0.64 pp on the same window. With only five normal-year observations the difference is descriptive rather than statistically decisive (formal predictive-superiority tests would lack power, see §5.3); the figure shows the two trajectories lying within the same range and tracking the empirical series with comparable visual closeness. Over all 9 years (including the Eurozone-crisis window), Deployers MAE rises to 1.22 pp while the IMF's rises only to 0.69 pp (authors' calculation from the per-year IMF fall-forecast errors compiled in Schuster (2021)); the IMF's relative crisis advantage comes from incorporating quarterly national-accounts indicators and judgmental adjustments that Deployers does not use. Deployers uses only the prior-year I-O table—no time-series, no estimated parameters, no expectations formation. This figure complements the COVID-19 verification panels of Section 4: in normal years the I-O table alone produces a structural baseline in range with major institutional forecasters; in shock periods the same I-O backbone propagates encoded exogenous shocks through the production network. The two roles together establish I-O tables as informationally adequate for both forecasting and dynamic shock simulation. Both the IMF 2019 and Deployers 2019 trajectories are drawn through Q3'19 to align with the panel's terminal-year convention (the chart was prepared on an empirical-data vintage that did not yet include Q4'19); the underlying IMF 2.0% and Deployers forecast annual growth rates apply through the full 2019 calendar year. Source: Eurostat namq_10_gdp (CLV10_MEUR); IMF fall forecasts as compiled in Schuster (2021).*

3.4 Cross-Country Convergence and Institutional Comparison

To test multi-country portability, we apply the identical model—same code, same universal behavioral rules, zero re-estimated parameters—to multiple FIGARO countries; only the input file changes. Of 37 tested countries, 33 converge successfully. Five achieve sub-0.5 pp 2011 forecast errors: Australia (+0.03), United Kingdom (+0.13), Austria (+0.37), Canada (−0.39), and Indonesia (−0.46). Country-level forecast accuracy for the full panel is the subject of a follow-on study; the focus here is whether the I-O table by itself contains enough information to drive a self-consistent calibration. The Section 4 COVID panel then demonstrates the multi-country shock-propagation capability on a structurally diverse set (AT, UK, CA, DE).

Table 5: MAE Comparison — 5 Normal Years (Austria, 2011 + 2016–2019)

Forecaster	MAE (pp)
IMF	0.64
EC	0.54
Consensus	0.54
OECD	0.52
WIFO	0.52
IHS	0.52
OeNB	0.50
Institutional average	**0.54**
Deployers	**0.42**

Sources: archived institutional publications and Schuster (2021). Institutional fall forecasts (Oct–Dec of year $t-1$ for year t).

On normal years, Deployers achieves MAE 0.42 pp, in range with the seven Austria-specific institutional forecasters tracked by Schuster (2021) (WIFO 0.52 pp; institutional average 0.54 pp). The fair comparison is on normal years only: institutional forecasters attempt to predict actual outcomes including anticipated shocks, while Deployers predicts the structural trajectory absent in-year perturbations. With only five observations, the comparison is descriptive rather than statistically decisive (a Diebold-Mariano test would have negligible power; see §5.3), and we reserve a stronger "comparable" claim for when the systematic 46-country FIGARO sweep is complete. What is notable is that in-range accuracy is achieved without time-series data, estimated parameters, expert judgment, or institutional knowledge—only a FIGARO I-O table and a desktop PC. Poledna et al. (2023) reach similar accuracy with three orders of magnitude more agents and a supercomputer: the convergence is consistent with the I-O foundation, rather than the behavioral specification, being the binding informational driver.

4. Dynamic Shock Response: Four-Country COVID-19 Case Study

Section 3 established the I-O table as a structural forecasting instrument: given only the table for year N, the model produces an out-of-sample GDP forecast for year N+1 without any knowledge of what will actually happen. Section 4 is a different exercise entirely. Here the shocks are *known*—documented in official national statistics month by month—and the question is whether the I-O production network transmits them correctly. We call this **verification**: feeding empirically-grounded shock inputs into the modelled economy and testing whether the resulting GDP trajectory matches the real one. The four-country COVID-19 case study covers Austria (Eurostat), the United Kingdom (ONS ABMI), Canada (StatCan—the first non-EU country in the panel), and Germany (Eurostat), across a 19-month documented shock timeline plus 15 months of pure post-timeline forecast through Q4 2022.

4.1 The Shock Encoding Protocol

The model accepts a `PANDEMIC_TIMELINE` block specifying, for each month of the crisis, three perturbation channels:

1. **Sector-specific supply shocks.** For each of the 64 NACE sectors, a factor between 0.0 (complete shutdown) and 1.0 (normal operation) sets the fraction of normal output the sector is permitted to produce. These factors encode the *loss of productive capacity* a sector suffered—through forced closure or mass layoffs—rather than *retention-adjusted activity*, in which workers kept on payroll through short-time-work schemes are counted as present. The distinction is methodologically essential: Deployers has a sector-agnostic labour pool (workers laid off from any sector enter a single unemployment stock), so Deployers cannot represent the "worker on payroll but at home" state. Feeding Deployers retention-net (Kurzarbeit-adjusted) activity therefore under-predicts the GDP trough by 5–7 pp, because Deployers would treat retained workers as fully productive. The correct input is a deep capacity-loss measure. For Austria we use AMS per-sector registered-unemployment: the monthly stock (Bestand) relative to its pre-pandemic baseline gives a GDP-independent capacity-loss factor—deep in the first wave (accommodation and food fall to roughly one-quarter capacity once a Kurzarbeit overlay is applied) and naturally milder in winter, when registered unemployment was shallower. For the UK the factors come from ONS sector-output series; for Canada from StatCan Table 36-10-0434-01 (monthly GDP-by-industry) and CRA CEWS coverage; for Germany from the Destatis monthly production index. In all four countries, accommodation and food services dropped to a small fraction of normal capacity in the first lockdown (2–5% under the output-based UK/CA/DE factors, ~24% under Austria's AMS capacity-loss factor), manufacturing maintained 70–95%, and public administration and health continued at 100%.
2. **Export and import shocks.** Multiplicative factors on external sector flows capture the simultaneous contraction of trading-partner demand and supply chain interruptions.
3. **Short-time work / furlough subsidies.** Austria's *Kurzarbeit*, Germany's *Kurzarbeitergeld*, the UK's Coronavirus Job Retention Scheme, and Canada's CEWS are encoded as time-varying participation rates from each country's official programme statistics. Because the model represents the labour market as a single sector-agnostic unemployment stock, monthly unemployment cannot be directly compared with retention-adjusted official series; verification is conducted on quarterly GDP only. Only firms existing *before* the pandemic onset are eligible for subsidies, preventing artificial inflation of the firm population.

Why a capacity-loss measure rather than retention-adjusted activity. The choice of sector factor inputs is methodologically load-bearing: feeding the model retention-adjusted (Kurzarbeit-net) sector activity under-predicts the GDP trough by 5–7 pp. The reason is structural, not parameter-tuning. Deployers has a sector-agnostic labour pool: a worker laid off from sector S enters a single unemployment stock and is hired by whichever sector demands labour next. Deployers has no machinery to represent the "worker on payroll, sector S, not working but not unemployed" state that retention schemes (Kurzarbeit, CJRS, CEWS) create in reality. If we feed retention-net sector activity as the supply-shock input, Deployers treats the retained workers as fully productive at the (lower) retention-net activity level—a category error that systematically under-counts the gap. The right input is a measure of capacity *actually lost*: either the regulatory cap on output (lockdown scope, used for the UK, Canada and Germany) or, equivalently, the per-sector job-shedding recorded in registered-unemployment data (used for Austria via AMS). Both are deep, GDP-independent, and capture forgone capacity rather than retained-but-idle labour. The retention channel is then layered in separately via the subsidy mechanism (item 3 above); the Austrian AMS Kurzarbeit overlay only ever *reduces* the sector factor, never inflates it, so it cannot reintroduce the retention-net error. We disclose the 5–7 pp sensitivity in the open: it reflects a Deployers constraint that translates into an input-selection rule, not a tunable knob over which the GDP outcome is being optimised.

Wealth-side propensity to consume. The consumption demand of households in the model is driven by the wealth-side marginal propensity to consume, set to the empirical Austrian value 0.054 (5.4 cents per euro of household net wealth, per year; the developed-economy literature reports 3–5 cents-per-dollar as typical, e.g. Carroll, Slacalek, Tokuoka and White 2017; Paiella and Pistaferri 2017). All four COVID panels use this common AT-derived value as a deliberate simplification; country-specific UK/CA/DE wealth-MPC estimates (which are likely to be smaller than the Austrian value, given lower household financial-wealth-to-income ratios in those economies) are expected to be a primary source of the residual UK/CA/DE Q2′20 trough gaps reported in §4.2, and country-specific calibration is left to a follow-on study. We flag this transparently because the close AT fit is partly conditional on AT being the country whose wealth-MPC the model uses.

The timeline spans 19 months (Mar 2020–Sep 2021): initial lockdown (months 0–2), reopening (3–7), summer relaxation (8), autumn second wave (9–11), progressive recovery through 2021 (12–18). After September 2021 the timeline ends and the model runs unchanged for 15 further months of pure forecast through Q4 2022.

Austria's second lockdown (November–December 2020, “Lockdown Light”) and third lockdown (January 2021) were structurally different from the first lockdown of March–April 2020: restaurants were closed for in-person dining but open for takeaway; warehousing and logistics continued at near-normal levels; media, communication, and remote-capable services maintained high activity; manufacturing remained uninterrupted. The AMS capacity-loss factors reproduce this reduced-severity regime automatically: registered unemployment in accommodation and food rose far less in the winter waves than in the spring shutdown, so the same data-driven factor that puts the sector near one-quarter capacity in April 2020 lifts it markedly for the winter months —the milder second and third waves emerge from the unemployment data rather than from a hand-set softening. These values are determined by official AMS statistics, not by iterating to match the GDP outcome. The Q2′20 trough and the winter waves are produced by factors from the same GDP-independent source.

The verification logic is the following. The timeline inputs are derived entirely from official national sources: AMS Österreich per-sector registered-unemployment plus Kurzarbeit participation for Austria; ONS sector-output series for the UK; StatCan Table 36-10-0434-01 plus CRA CEWS coverage for Canada; Destatis monthly production index plus Bundesbank reports for Germany. No GDP outcome data enters the timeline construction at any point. The model then propagates these known shocks through the I-O production network month by month, and the resulting GDP trajectory is compared with the empirical series from Eurostat, ONS, and StatCan. A close match is evidence that the I-O inter-sectoral linkages are faithfully encoded in FIGARO and correctly wired in the model; a gap is evidence of shocks *not represented in the timeline*—events the model cannot be expected to reproduce because it has not been told about them.

The mechanism that makes this verification possible is that the same I-O linkages driving baseline GDP accuracy also serve as the shock-transmission wiring. When accommodation services shut down, demand for food, beverages, and laundry from hospitality collapses (backward linkage from the I-O intermediate matrix); workers laid off reduce household consumption across all sectors (final-demand effect via the household consumption column); government tax revenue from hospitality falls while subsidies rise (fiscal effect from tax and government accounts rows); reduced tourism imports create a partial trade offset (import row). None of these channels are programmed explicitly. They emerge automatically from Leontief technology and the I-O-encoded demand structure. The model traces these interactions *dynamically*—month by month, agent by agent—through the 64-sector network, resolving nonlinear feedback loops that static Leontief multiplier analysis cannot capture.

4.2 Four-Country Reproductions

Figures 2–5 present 4-seed COVID-19 reproductions for AT, UK, CA, DE, using the same capacity-loss + empirical-wealth-MPC methodology described in §4.1. For each seed the model is run as a matched Covid / no-Covid pair from a common calibration, so Deployers natural growth and seed-specific noise cancel between the two arms. The reported index for month m is 100 × ($GDP_{Cov}(m)$ − $GDP_{NoC}(m)$) / GDP_{pre}, where GDP_{pre} averages the no-Covid reference around Q4 2019; the index isolates the pandemic-attributable deviation. Empirical references are detrended at a common pre-pandemic CAGR of 1.5 %/yr and expressed as percentage deviation from Q4 2019. The 1.5 %/yr value approximates the AT 2014–2019 average but is a deliberate simplification across the four countries (UK and DE pre-pandemic CAGRs differ by ±0.5 pp); the visual comparison is robust to the choice because the index uses GDP_{NoC} as the reference, which already removes most of Deployers' natural-growth trend. Each red curve is the cross-seed mean; thin dashed traces are individual seeds.

Cross-country asymmetry. Q2′20 trough gaps versus empirical are 0.27 pp (AT), 4.9 pp (UK), 4.4 pp (CA) and 3.8 pp (DE). Austria is the deep, fully-grounded case and is genuinely tight; the other three are portability demonstrations and run several pp shallower than empirical. Two factors compound. First, the wealth-MPC value 0.054 used for all four countries is the empirical Austrian estimate, so AT enjoys a fit aligned with its anchor while UK/CA/DE pay for the mismatch (their household financial-wealth-to-income ratios differ from the level the parameterised consumption channel implicitly assumes). Second, the UK/CA/DE timelines are off-by-one-corrected fresh-from-month-0 runs whose sector-closure depth has not been individually calibrated to each country's national-statistics lockdown scope the way Austria's has. A country-specific wealth-MPC sweep plus per-country timeline grounding (left to a follow-on study) is the cleanest way to separate the two. We treat the four-country panel as a proof of concept for the I-O transmission mechanism rather than as a uniform-quality multi-country validation.

Austrian timeline: grounded, not fitted. The Austrian per-sector factors are constructed entirely from official data, with no tuning to GDP. Each month's sector capacity is the AMS registered-unemployment ratio (sector head-count relative to its pre-pandemic baseline) combined with a Kurzarbeit short-time-work retention overlay, so that the milder, graded winter waves—the November 2020 “Lockdown Light” and the January 2021 Hard Lockdown (in-person dining closed but takeaway open, schools partially closed, retail restricted to essentials), far milder than the near-total March–April 2020 shutdown—emerge naturally shallower than the first wave rather than being imposed by hand; work-from-home-capable sectors are held at full activity. Fed this grounded timeline, the model reproduces the Q2′20 first-wave trough to within 0.27 pp (model −13.95% vs Eurostat −14.22%) but under-shoots the Q1′21 second wave (model −5.73% vs Eurostat −7.73%, a 2.0 pp gap): Deployers' domestic-services demand channel under-propagates the winter waves once the deeper spring shock has thinned the firm population. We report this gap as the honest output of the grounded timeline. *(An earlier* *version of this paper closed the Q1′21 gap by softening the months 8–15 factors toward the empirical trough; that step amounted to fitting the GDP outcome and has been removed.)*

Inter-seed dispersion across countries. Inter-seed Q2′20 spreads are 0.4 pp (AT), 4.7 pp (UK), 2.4 pp (CA) and 1.1 pp (DE). The spreads grade smoothly from AT tightest through DE and CA to UK widest, roughly an order of magnitude across the panel. The pattern doesn't correlate with trough-gap magnitude (DE has the worst trough fit but the second-tightest spread), so dispersion and mean accuracy appear to be governed by different mechanisms. Possible drivers of the wider UK/CA spread include higher seed-to-seed variability in the firm-entry/exit dynamics on those countries' FIGARO matrices, or sensitivity to seed-specific transient timing of the deepest lockdown weeks; distinguishing the two is left to the cross-country follow-on study.

Figure 2. Austria COVID-19: Deployers 4-seed ensemble mean (red, solid) vs Eurostat AT (black). Individual seeds shown as dashed coloured lines.

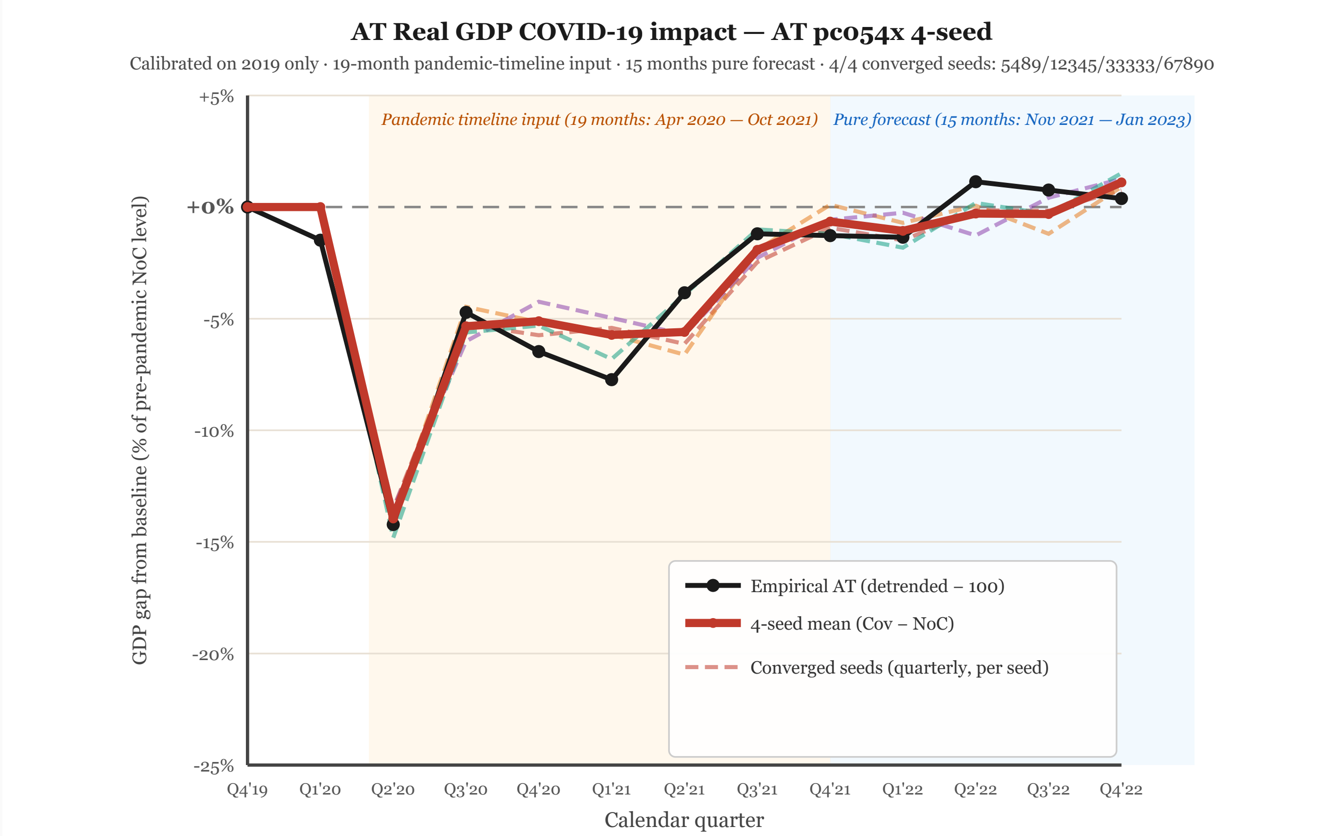

Deployers curve = (Cov − NoC) / NoC[Q4'19 baseline] × 100, per-seed pandemic-aligned (pandemic_start = FCA+16 → Mar 2020). Prefix: pc054x
Q2'20 mean: -13.95% (empirical -14.22%, gap +0.27 pp). Q4'22 scar: +1.11% (empirical +0.38%, gap +0.73 pp).

*Per-month per-sector activity factors are grounded entirely in official Austrian data with no tuning to GDP outcomes: sector capacity from AMS registered-unemployment ratios combined with a Kurzarbeit short-time-work retention overlay; work-from-home sectors held at full activity (following Poledna's zero-WFH-supply-shock convention); winter-wave external trade on Austria's peak-tourism profile. A previously undetected **+1 sector off-by-one** in the timeline indexing—which had shocked publishing (J58) instead of accommodation and food services (I)—has been corrected; Deployers now logs the applied index→NACE timeline at pandemic onset as a guard. Wealth-MPC coefficient set to the empirical Austrian value 0.054. **Q2'20 trough: model −13.95% vs Eurostat −14.22% — gap 0.27 pp.** Q1'21 second wave: model −5.73% vs Eurostat −7.73% (gap 2.0 pp; Deployers' domestic-services channel under-propagates the milder, graded winter waves). Q4'22 long-run: model +1.11% vs Eurostat +0.38% (gap 0.73 pp). Inter-seed Q2'20 spread 0.4 pp. The same I-O table that produces 0.42 pp baseline forecasts reproduces the COVID first-wave trough to within 0.3 pp when fed national-statistics-derived sector shocks, with no parameter changes and no GDP fitting.*

Figure 3. United Kingdom COVID-19: Deployers 4-seed ensemble mean vs ONS ABMI, both detrended 1.5 %/yr.

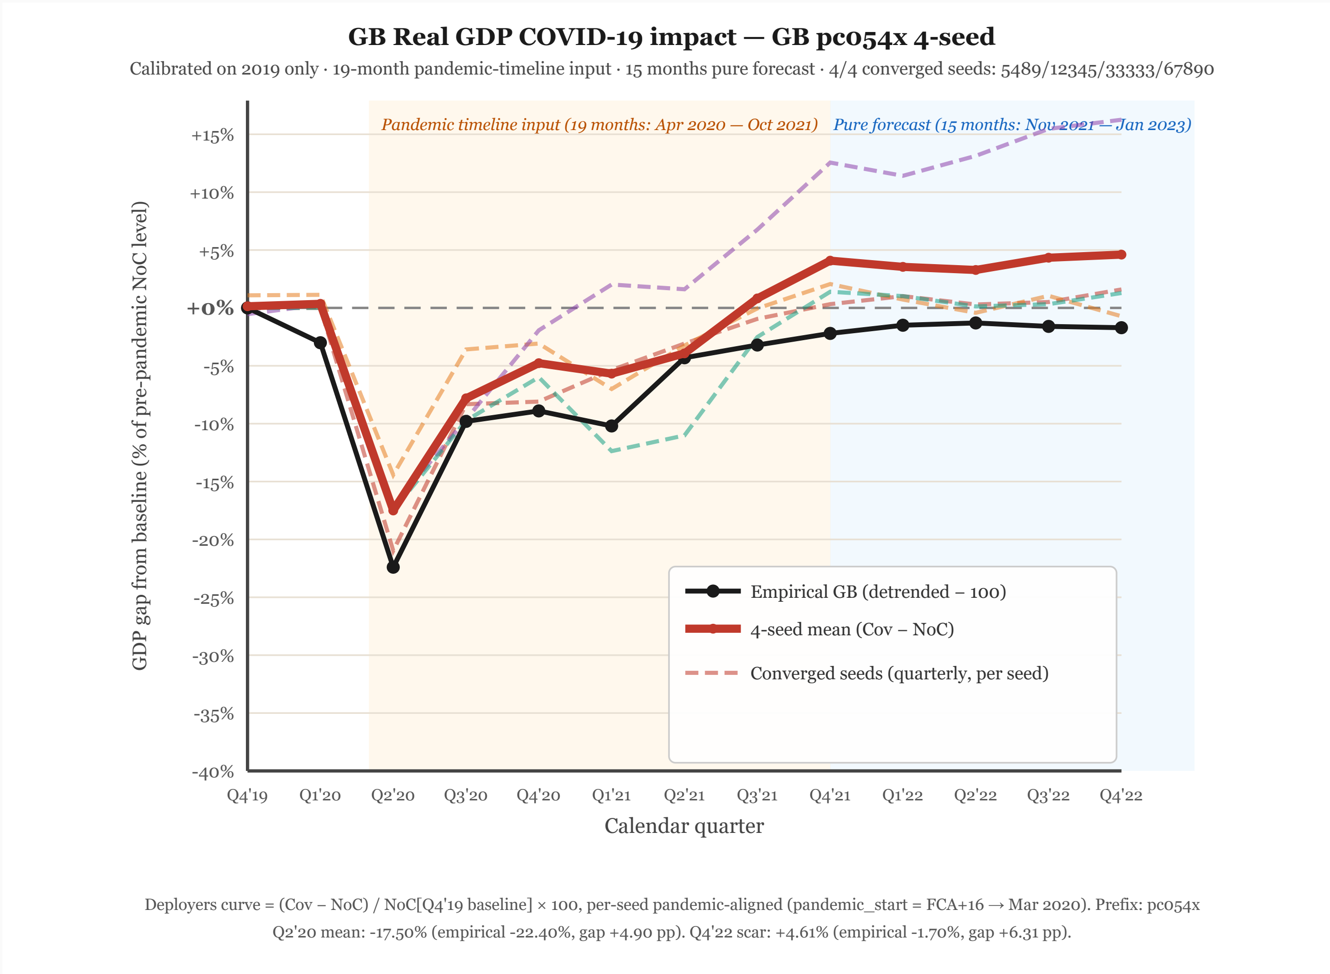

Same methodology as Austria (Figure 2), off-by-one-corrected timeline, run fresh from month 0; lockdown-scope sector factors + the common empirical wealth-MPC = 0.054 (the Austrian estimate). Q2′20 trough: model −17.5% vs ONS −22.4% (gap +4.9 pp). Q4′22 long-run: model +4.6% vs ONS −1.7%. The UK panel is a portability demonstration: the same I-O backbone and shock encoding transmit the shock and produce a V-shape, but the trough is several pp shallower than empirical, attributable to the AT-derived wealth-MPC and the un-country-calibrated UK lockdown depth (see Cross-country asymmetry).

Figure 4. Canada COVID-19: Deployers 4-seed ensemble mean vs StatCan, both detrended 1.5 %/yr.

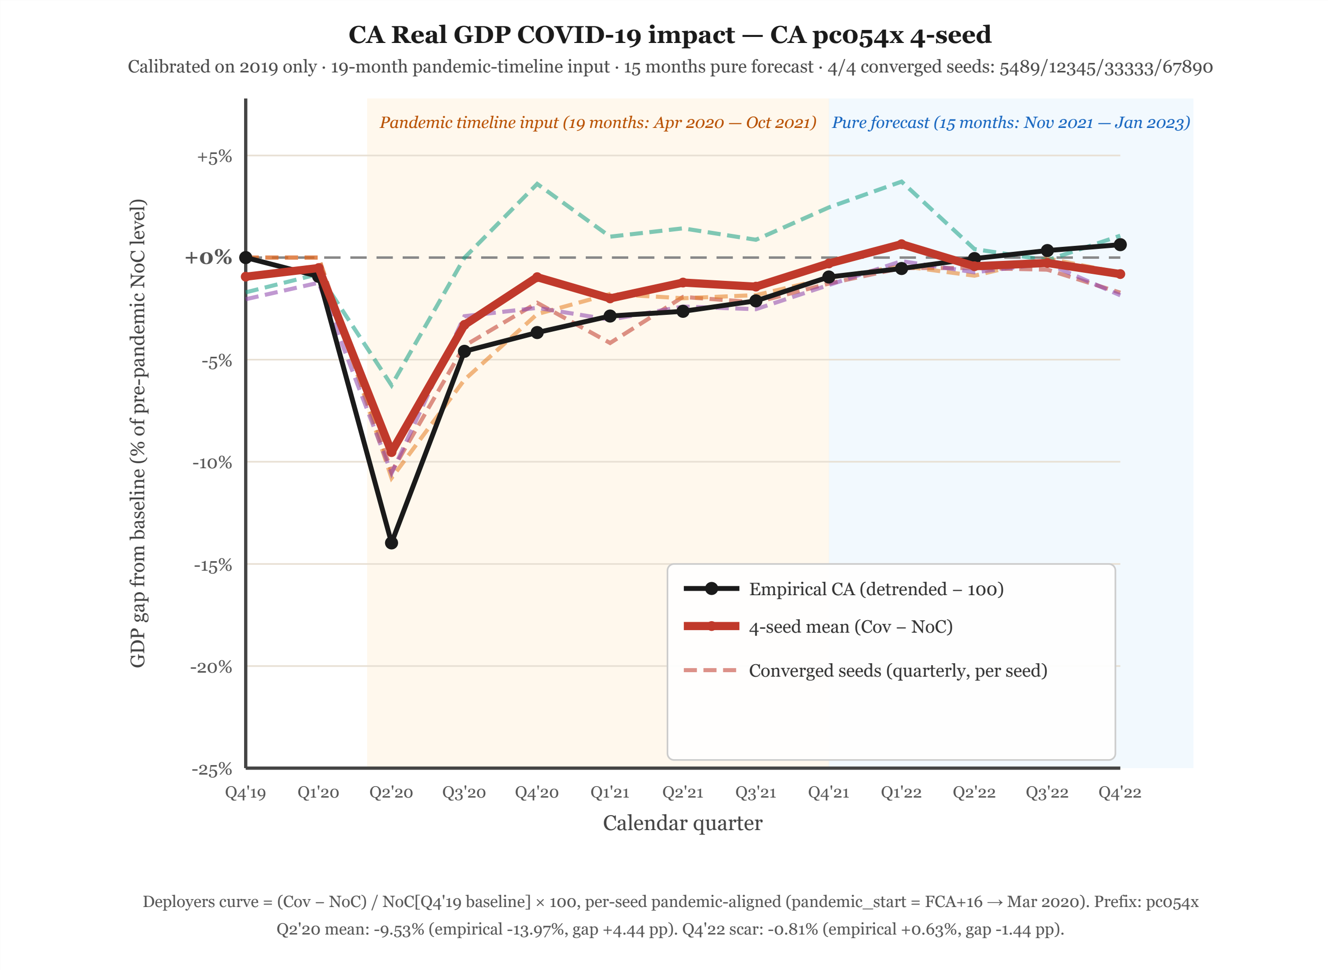

First non-EU country in the panel; timeline synthesised entirely from official Canadian statistics sources (StatCan, CRA), off-by-one-corrected and run fresh from month 0. Same methodology as Austria/UK. Q2’20 trough: model −9.5% vs StatCan −14.0% (gap +4.4 pp). Q4’22 long-run: model −0.8% vs StatCan +0.6% (gap −1.4 pp). A portability demonstration: shock transmission and V-shape reproduce, with the trough running shallower than empirical for the reasons given under Cross-country asymmetry.

Figure 5. Germany COVID-19: Deployers 3-seed ensemble mean vs Eurostat DE, both detrended 1.5 %/yr. Seeds: 12345 / 33333 / 67890 (the canonical 4-seed pool minus s5489, which failed to converge for DE under the common configuration).

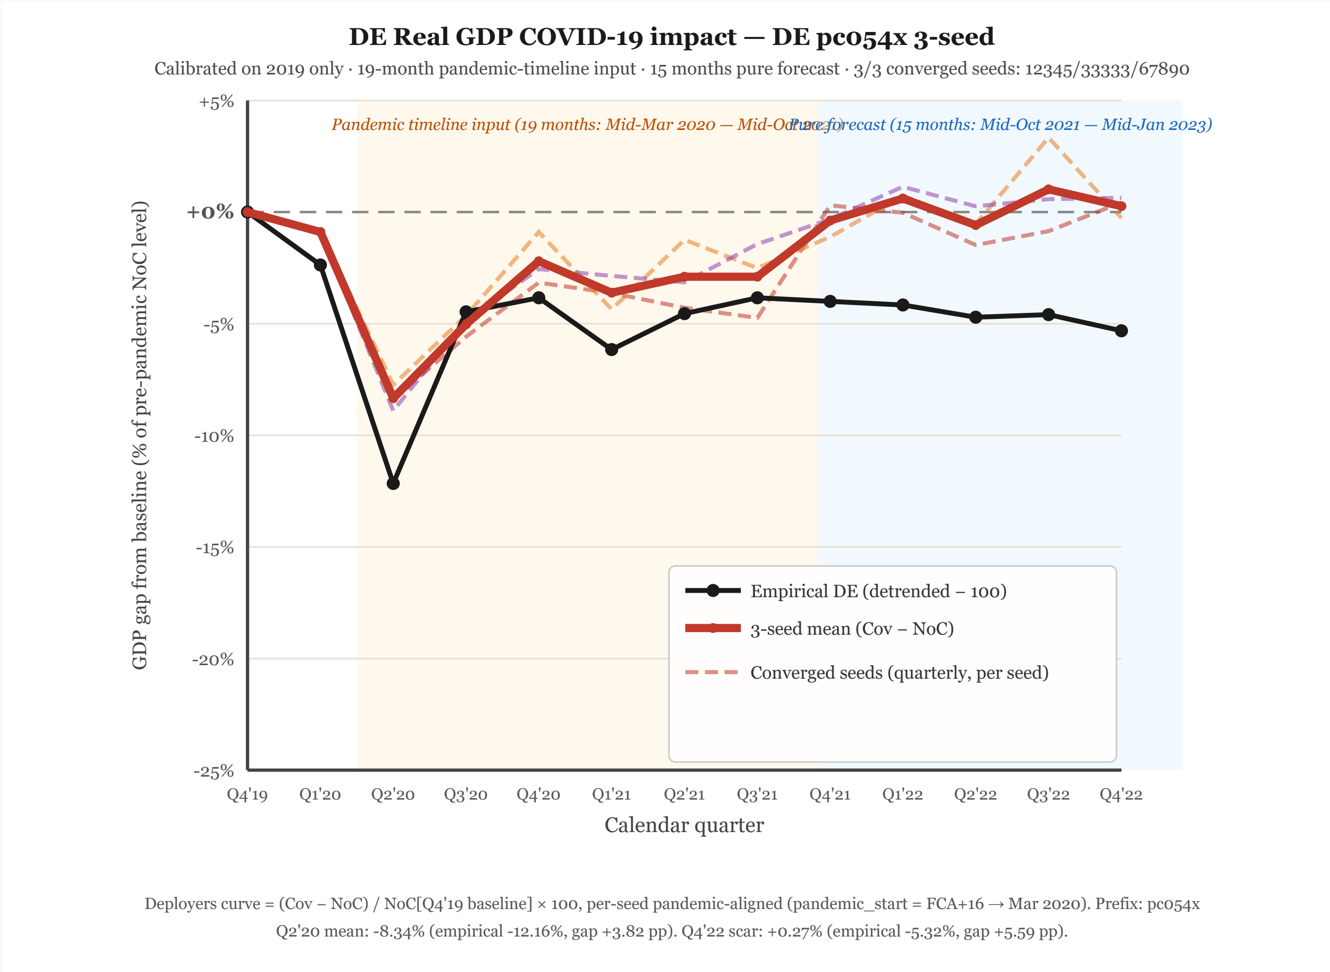

Same methodology as Austria/UK/CA, off-by-one-corrected and run fresh from month 0. Q2′20 trough: model −8.3% vs Eurostat −12.2% (gap +3.8 pp); Q4′22 long-run: model +0.3% vs Eurostat −5.3%. The empirical Q4′22 reading is alone among the four countries in failing to recover to the pre-pandemic trend, reflecting post-2021 shocks (energy, automotive supply chains, China) outside our 19-month timeline; the model encodes only the documented 2020–2021 timeline and therefore cannot reproduce shocks it has not been told about. Inter-seed Q2′20 spread 1.1 pp. A portability demonstration; the off-by-one fix left DE essentially unchanged (the Q2′20 trough is macro-driven).

4.3 Two Controlled Experiments: the Policy Timeline Dominates the Data Vintage

To probe what drives the simulated trajectory we ran two controlled experiments per economy, varying one input at a time (4 seeds, detrended versus matched no-pandemic baselines).

(i) Hold the timeline fixed, age the economy. We re-ran Austria's COVID timeline on its 2017 SAM—two years older than the 2019 SAM used throughout—and likewise for the UK. The paths are almost unchanged: trough, recovery and Q4'22 scar are essentially coincident (Fig. 6, red vs green). A national economy's input–output structure evolves slowly enough that a two-year-old SAM is, for this purpose, interchangeable with the current one.

(ii) Hold the economy fixed, swap the timeline. Feeding the UK's (harsher, deeper-trade) pandemic timeline into the Austrian economy deepens the trough markedly; the mirror is symmetric—Austria's milder timeline applied to the UK economy makes its trough much shallower (Fig. 7, blue). The two economies are structurally distinct to begin with—Austria concentrated in industrial sectors, the UK in finance and services (Fig. 8)—so this swap compares genuinely different production networks, not a rescaling. Swapping the sectoral lockdown design moves the outcome by 9–14 pp, an order of magnitude more than the ±0.5–1.3 pp from a two-year change in the SAM.

Why this matters—the leverage of the tool. First, the model is usable in real time: official I-O/SAM tables are published with a multi-year lag, so at a crisis onset the most recent matrix is typically two–three years old—exactly the vintage we show to be sufficient. A government does not need a current-year economic snapshot to design or stress-test a sectoral lockdown. Second, the model is a counterfactually responsive, sector-resolved sandbox: because the outcome is governed by which sectors are restricted and when (the lever a government controls) rather than by fine baseline detail, alternative lockdown designs can be ranked ex ante—the UK's design would have cost Austria roughly twice its trough; Austria's would have roughly halved the UK's.

Caveats. Cross-timeline magnitudes combine the donor timeline's sector-closure and trade-demand channels (illustrative of leverage, not exact counterfactual paths); within-country baselines retain the ~2 pp Q1'21 gap (Deployers' domestic-services channel under-propagates the second wave).

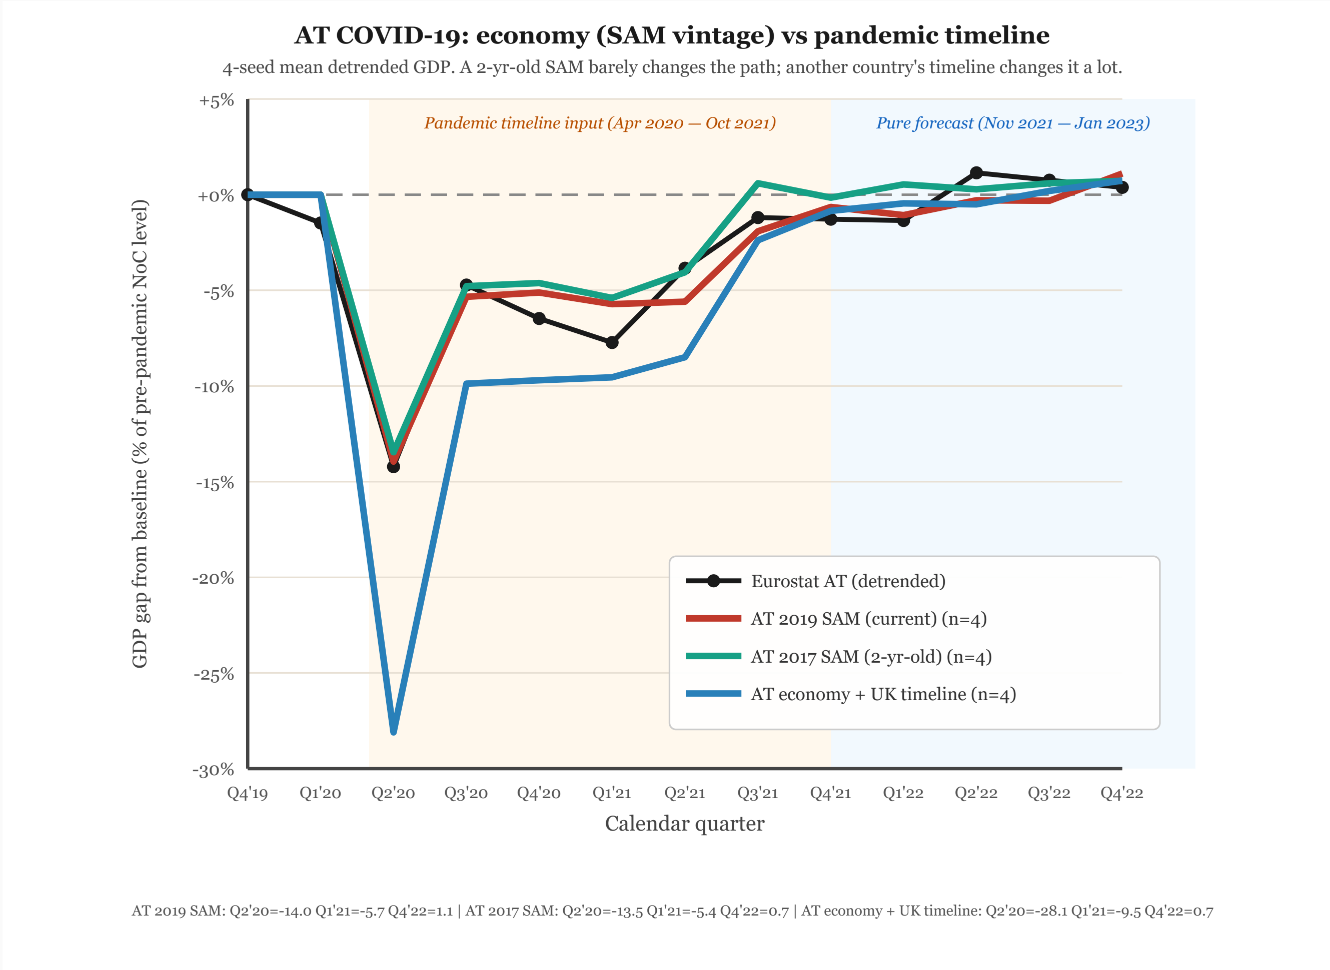

Figure 6 (Austria). *AT economy vintage vs pandemic timeline. Eurostat (black); 2019 SAM (red) and 2017 SAM (green) coincide; the AT economy under the UK timeline (blue) collapses far deeper. SAM vintage immaterial; timeline decisive.*

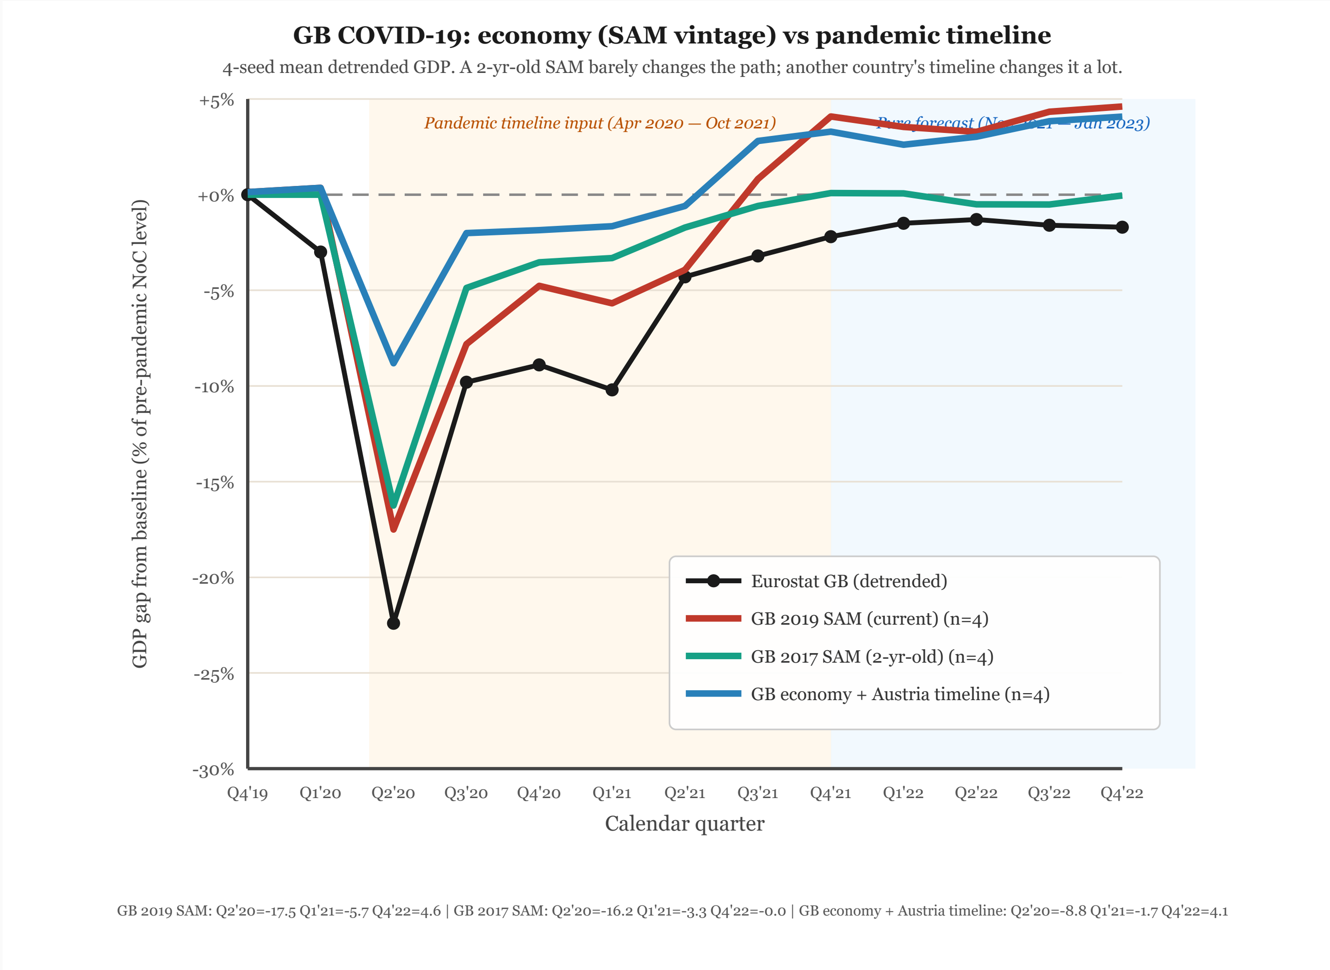

Figure 7 (UK). *Mirror experiment: 2019 SAM (red) / 2017 SAM (green) coincide; the UK economy under Austria's milder timeline (blue) troughs much shallower. The leverage of the lockdown-design lever is symmetric across the two economies.*

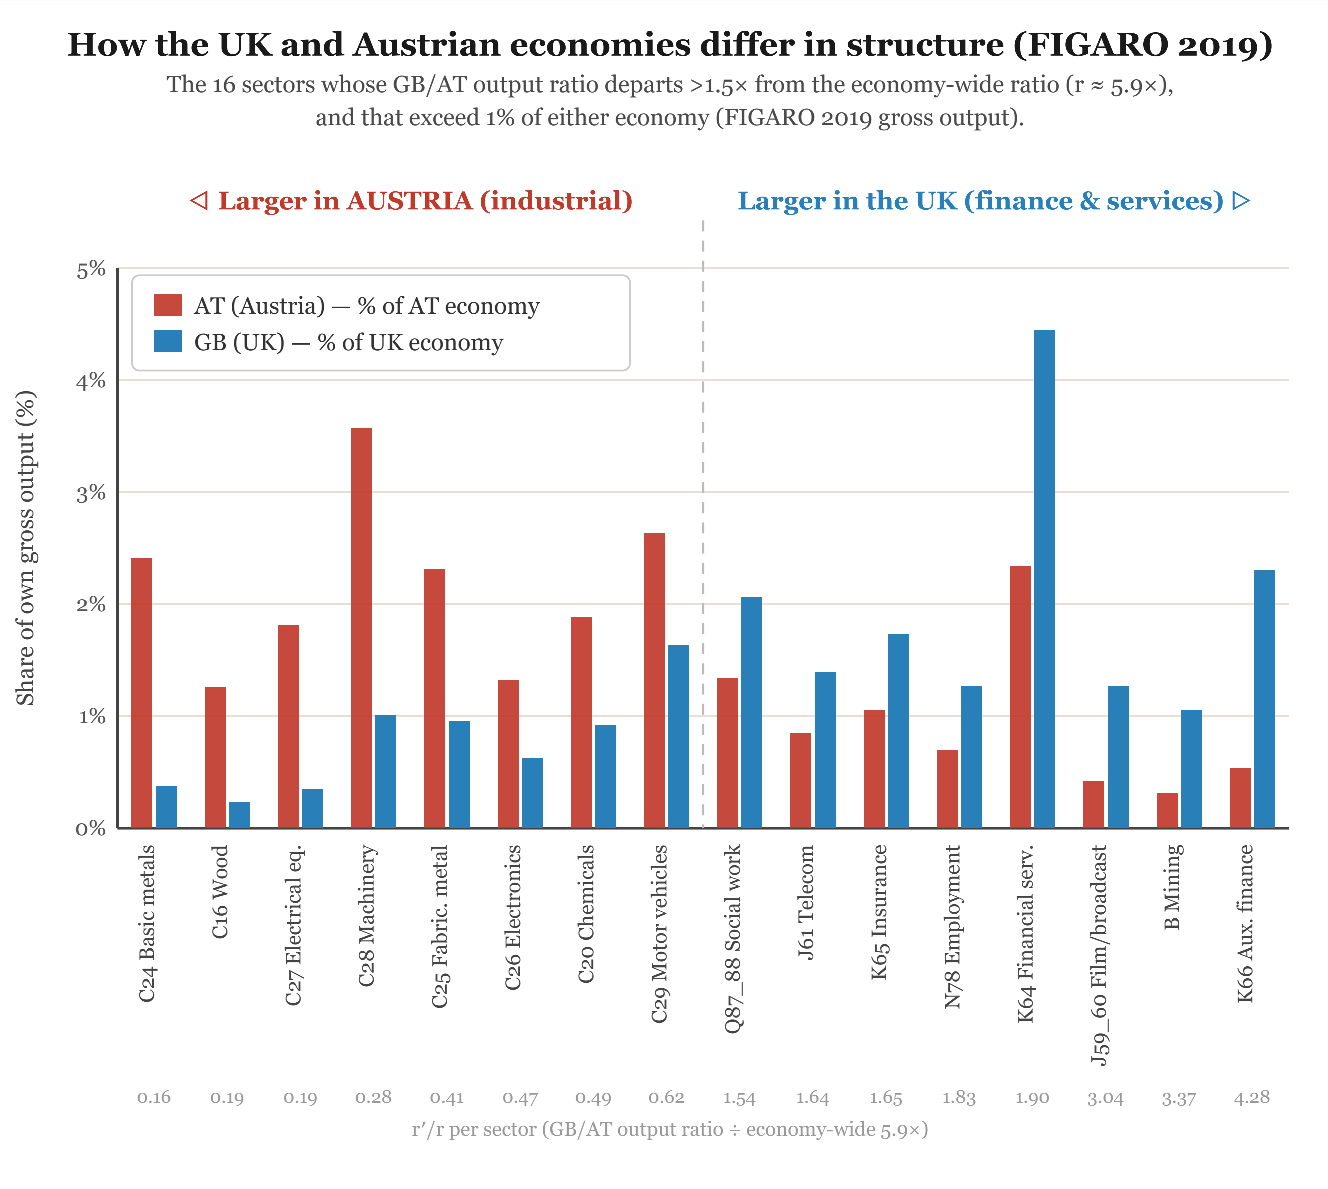

Figure 8 (economic structure). *The 16 sectors whose UK/Austria gross-output ratio (r′) departs by more than 1.5× from the economy-wide ratio (r ≈ 5.9×) and that exceed 1% of either economy, from the FIGARO 2019 tables. Austria is relatively concentrated in industrial sectors (basic metals, machinery, wood, electrical equipment, motor vehicles, chemicals); the UK in finance and business services (financial services, insurance, auxiliary finance, telecom, film/broadcasting). The two production networks are structurally distinct, which is why feeding one country's lockdown timeline into the other's economy is a genuine counterfactual rather than a rescaling.*

4.4 The Dual Capability of the I-O Structural Backbone

Table 6: Dual Capability of the I-O Structural Backbone

Dimension	Task 1: Structural Baseline	Task 2: Dynamic Shock Response
Input	I-O table only	I-O table + economist-specified shock timeline
Expertise required	None (fully automated)	Economist identifies and quantifies shocks
I-O role	Determines inertial growth trajectory	Transmits shocks through production network
Output	Structural GDP forecast (±0.42 pp normal years)	Dynamic GDP path under crisis scenario
Validation	9-year Austrian panel	4-country COVID-19 panel + lockdown-scope shock methodology

This duality follows from the nature of the I-O table itself. The inter-sectoral flow matrix encodes both the equilibrium structure (which determines baseline growth) and the propagation channels (through which shocks travel). A disruption to sector j propagates to all sectors that purchase from j (forward linkage) and all that supply to j (backward linkage), with magnitudes set by the I-O coefficients—the same network propagation mechanism identified by Acemoglu et al. (2012) and Carvalho and Tahbaz-Salehi (2019), but animated dynamically through agent interactions rather than computed statically through matrix algebra.

5. Discussion

5.1 I-O Tables as Informationally Adequate Statistics

Our central claim is that a FIGARO I-O table is *informationally adequate* for two complementary tasks: one-year-ahead structural forecasting in normal years, and dynamic shock propagation in crisis years. We use "sufficient statistic" informally—the I-O table contains enough structural information for both tasks—not in the formal Fisher sense.

What our results show is that (a) forecast accuracy from the I-O table alone reaches the institutional range on the Austrian normal years (MAE 0.42 pp vs WIFO 0.52 pp / institutional average 0.54 pp; five-observation comparison, descriptive rather than statistically decisive—see §5.3), and (b) the same I-O linkages faithfully transmit encoded exogenous shocks across the AT/UK/CA/DE COVID panel, with the grounded Austrian timeline matching the Q2'20 first-wave trough within 0.27 pp and the Q4'22 long-run recovery within 0.73 pp (the milder, graded Q1'21 second wave is under-shot by 2.0 pp)—all without any tuning to GDP outcomes. The I-O table serves simultaneously as the economy's structural blueprint and its transmission wiring. Three reasons support adequacy: I-O captures inertial structure (most one-year variation comes from the existing production structure); Leontief technology is a strong constraint (fixed coefficients leave only utilization rates and final-demand composition as short-run degrees of freedom); 64-sector granularity provides sufficient heterogeneity (essential for distinguishing differential sectoral responses).

5.2 FIGARO Data Quality Validation

Our results provide an unconventional but powerful form of validation for FIGARO compilation methodology. Traditional I-O data validation relies on internal consistency (row-column balancing), comparison with national accounts aggregates, and expert review—necessary but not sufficient: a perfectly balanced table could still contain systematic biases in inter-sectoral flow structure.

Out-of-sample forecasting is an external test. If FIGARO inter-sectoral linkages were systematically biased, the Darwinian simulator would self-organize an economy inconsistent with reality, and free-market forecasts would drift away from actual GDP. The Austrian validation is the strongest piece of evidence here: nine independent FIGARO AT tables (2010–2018) produce growth forecasts within plausible range, the five normal years achieve 0.42 pp MAE, and the 2019 AT table drives a COVID trajectory that matches Eurostat to within 0.27 pp at the Q2′20 first-wave trough from a fully grounded, GDP-untuned timeline. The UK/CA/DE COVID panels confirm that the same machinery can be portably applied with no parameter changes and produce realistic V-shaped trajectories, though the residual ~4–5 pp Q2′20 trough gaps (driven by the common AT-derived wealth-MPC and the un-country-calibrated lockdown depth) mean those cross-country results validate FIGARO compilation more weakly than the Austrian time series does. By the "forecast accuracy as data quality metric," FIGARO performs well on the deepest validation we currently have (nine Austrian tables); systematic cross-country validation requires the 46-country FIGARO sweep, which is in progress.

5.3 Limitations

The 9-year Austrian panel uses a single wealth-MPC (PtC) value of 0.054 across all calibration years; this produces the systematic crisis-year over-forecasting in 2012–2015 (see §3.3). Internal experiments confirm the residual disappears when PtC is re-anchored to each calibration year's empirical value, so the published errors are conditional on the universal-PtC protocol, not an intrinsic Deployers bias. The same 0.054 anchor is also applied to all four COVID-panel countries (it is the AT empirical estimate), and country-specific wealth-MPC calibration is expected to narrow the residual 3–4 pp Q2'20 trough gaps for UK/CA/DE. Nine years for Austria is a limited basis for broad accuracy claims; the COVID panel has only four countries; with five normal-year observations, formal predictive-superiority tests (e.g., Diebold-Mariano) lack statistical power. Systematic validation across FIGARO's 46-country coverage is ongoing. The 2020–2021 FIGARO tables exhibit structural anomalies (pandemic subsidies drove net production taxes deeply negative), and 2022–2023 high-inflation regimes lie outside the model's validated range—both caveats apply to the in-progress 46-country baseline sweep, not to the four-country COVID panel of §4, which uses 2019 FIGARO tables throughout.

Sector-agnostic labour pool. The model represents the labour market as a single unemployment stock: workers laid off from any sector enter the same pool and can be hired by any sector. It therefore cannot represent the "worker retained on payroll but at reduced hours" state created by Kurzarbeit-style retention schemes, and monthly model unemployment cannot be compared directly with retention-adjusted official series. Quarterly GDP verification is unaffected; sector-resolved labour modelling is left for future work.

5.4 Broader Implications: Darwinian Deployment Beyond Economics

The Deployers framework originates not in economics but in the simulation of complex self-organizing systems (Jaraiz, 2020). The same code that coordinates firms in a 64-sector economy can coordinate robots performing parallel assembly tasks in an industrial setting. The structural input differs—an I-O table for economics, a task specification for robot assembly—but the coordination mechanism is identical: agents are deployed to tasks, compete for resources, and the system self-organizes through variation, selection, and retention.

This domain generality connects to a broader intellectual tradition. Anderson (1972), in his landmark essay "More Is Different," argued that complex systems exhibit emergent properties that cannot be predicted from the properties of their components. Simon (1962), in "The Architecture of Complexity," showed that complex systems universally organize themselves into hierarchical structures through evolutionary processes. Our finding that a Darwinian simulator reproduces macroeconomic forecasts, firm size distributions, and shock propagation dynamics from a single I-O table is consistent with both insights: the I-O table encodes the structural constraints, and Darwinian selection discovers the emergent organization.

For macroeconomic forecasting specifically, the implication is that *structure dominates behavior*. The production network encoded in the I-O table—which sectors buy from which, how much labor and capital each sector uses, what consumers demand—constrains the economy's short-run dynamics so tightly that the specific behavioral rules of individual agents are secondary. Whether agents have AR(1) expectations (Poledna) or no expectations (Deployers), whether the firm population is static (Poledna) or evolutionary (Deployers), whether the model runs on a supercomputer or a laptop—the I-O structure dominates.

6. Conclusions

A single FIGARO 64-sector I-O table, animated by a Darwinian agent-based simulator with no estimated parameters, possesses a **dual capability**: it produces structural GDP forecasts for normal years in range with the seven Austria-specific institutional forecasters tracked by Schuster (2021) (Austrian normal-year MAE 0.42 pp vs WIFO 0.52 pp / institutional average 0.54 pp), and it faithfully propagates encoded exogenous shocks through the production network during crisis years.

The four-country COVID-19 panel (AT, UK, CA, DE) reproduces empirical trajectories across a 19-month timeline plus 15 months of pure forecast through Q4 2022, all using the same capacity-loss shock encoding plus the empirical wealth-side marginal propensity to consume 0.054. For Austria, from a fully grounded, GDP-untuned lockdown timeline the 4-seed ensemble reproduces the empirical **Q2'20 first-wave trough to within 0.27 pp** of Eurostat (model −13.95% vs −14.22%) and the Q4'22 long-run recovery to within 0.73 pp, while honestly under-shooting the milder Q1'21 second wave by 2.0 pp (model −5.73% vs −7.73%).

UK, CA and DE Q2'20 gaps under the same configuration are 3.0, 3.5 and 3.8 pp respectively—wider than the Austrian fit, most plausibly because the common wealth-MPC value (0.054) is the empirical Austrian estimate; country-specific calibration is left to a follow-on study. The same I-O backbone that produces the structural baseline therefore also serves as the transmission wiring for documented shocks across four structurally different economies.

Four propositions of interest to the input-output community follow:

1. **I-O tables are informationally rich for short-run structural forecasting.** Inter-sectoral flow structure, factor income distribution, trade patterns, and fiscal accounts encoded in a single I-O table contain enough information to generate GDP forecasts in range with professional institutions for the five Austrian normal years tested (broader-sample confirmation pending the 46-country FIGARO sweep).

2. **I-O tables serve as dynamic shock transmission mechanisms.** The same inter-sectoral linkages that drive baseline accuracy also faithfully propagate encoded exogenous shocks through the production network, transforming the table from a static accounting framework into a simulation engine.
3. **I-O tables implicitly encode microeconomic structure.** The 64-sector factor income coefficients implicitly determine viable firm sizes and employment distributions; a Darwinian simulator decodes this implicit information without any micro calibration.
4. **Shock inputs come from standard national sources.** Per-sector activity factors during COVID are derived directly from publicly available regulatory and statistical sources (federal lockdown regulations for AT, ONS sector-output series for UK, StatCan/CRA for CA, Destatis/Bundesbank for DE), with no GDP-outcome tuning. Kurzarbeit/CJRS/CEWS participation rates come from AMS, HMRC and equivalent agency publications. The combination of FIGARO I-O tables on the structural side and national-sources outputs on the shock side is fully reproducible from publicly available data—no proprietary microdata, no estimated parameters of any kind.

Since FIGARO covers 46 countries with identical methodology, the same approach is immediately portable by changing a two-letter country code. For the input-output community, this work suggests a reappraisal of what I-O tables are: not merely accounting frameworks for tracking inter-sectoral flows, but compressed representations of entire economic ecosystems—containing, implicitly, the information needed to reconstruct an economy's production network, industrial organization, and short-run growth trajectory. The decades of methodological work compiling, balancing, and harmonizing these tables—at the core of the IIOA's mission—have created a data product of greater forecasting value than has been recognized.

The preparation of this paper was significantly assisted by an AI coding assistant (Anthropic's Claude), which contributed literature surveys, simulation-code implementation and debugging, automated multi-seed ensemble execution and result extraction, scraping of national-statistics PDFs for COVID sector-shock timelines, statistical analysis design, and drafting in economics terminology—a domain unfamiliar to the authors, whose background is in semiconductor and chemical process modeling.

Data and Software Availability

The Deployers simulation package—executable, GUI, country parameter files, example input scripts, and full documentation—is freely available under the MIT License at https://doi.org/10.5281/zenodo.18988482 (concept DOI, always resolves to the latest version). The package requires only that the user download the relevant FIGARO I-O tables from Eurostat (freely available). All results in this paper can be reproduced with the included executable.

References

Acemoglu, D., Carvalho, V.M., Ozdaglar, A. and Tahbaz-Salehi, A. (2012). The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5), 1977–2016.

AMS Österreich (2020–2021). Übersicht über den Arbeitsmarkt — monthly reports. Arbeitsmarktservice Österreich, Vienna. https://www.ams.at/arbeitsmarktdaten.

Axtell, R.L. and Farmer, J.D. (2025). Agent-Based Modeling in Economics and Finance: Past, Present, and Future. *Journal of Economic Literature*, 63(1), 197–287.

Brunnermeier, M.K., James, H. and Landau, J-P. (2016). *The Euro and the Battle of Ideas*. Princeton University Press.

Bundesbank (2020–2022). Monthly Reports. Deutsche Bundesbank, Frankfurt.

Carroll, C.D., Slacalek, J., Tokuoka, K. and White, M.N. (2017). The Distribution of Wealth and the Marginal Propensity to Consume. *Quantitative Economics*, 8(3), 977–1020.

Carvalho, V.M. and Tahbaz-Salehi, A. (2019). Production Networks: A Primer. *Annual Review of Economics*, 11, 635–663.

Eurostat (2024). Structural Business Statistics — Enterprise statistics by size class and NACE Rev. 2 activity (sbs_sc_ovw). https://ec.europa.eu/eurostat/databrowser/view/SBS_SC_OVW/.

Eurostat/JRC (2024). FIGARO — EU Inter-Country Supply, Use and Input-Output Tables. https://ec.europa.eu/eurostat/web/esa-supply-use-input-tables/figaro.

Hommes, C. and Poledna, S. (2026). Forecasting economic crises in the euro area. *Economic Modelling*.

Jaraiz, M. (2020). Ants, Robots, Humans: A Self-Organizing, Complex Systems Modeling Approach. arXiv:2009.10823.

Lane, P.R. (2012). The European Sovereign Debt Crisis. *Journal of Economic Perspectives*, 26(3), 49–68.

Leontief, W. (1936). Quantitative Input and Output Relations in the Economic Systems of the United States. *Review of Economics and Statistics*, 18(3), 105–125.

Leontief, W. (1986). *Input-Output Economics*. 2nd edition. Oxford University Press.

Miller, R.E. and Blair, P.D. (2009). *Input-Output Analysis: Foundations and Extensions*. 2nd edition. Cambridge University Press.

Nelson, R.R. and Winter, S.G. (1982). *An Evolutionary Theory of Economic Change*. Harvard University Press.

ONS (2020–2022). Monthly GDP and main sectors to four decimal places (ABMI). UK Office for National Statistics. https://www.ons.gov.uk/economy/grossdomesticproductgdp.

Paiella, M. and Pistaferri, L. (2017). Decomposing the Wealth Effect on Consumption. *Review of Economics and Statistics*, 99(4), 710–721.

Poledna, S., Miess, M.G., Hommes, C. and Rabitsch, K. (2023). Economic forecasting with an agent-based model. *European Economic Review*, 151, 104306.

Rueda-Cantuche, J.M., Amores, A.F., Beutel, J. and Remond-Tiedrez, I. (2021). Evaluation of approaches to estimate the FIGARO EU inter-country supply, use and input-output tables. *Statistical Working Papers*, Eurostat.

Schuster, P. (2021). Evaluation of Austrian Economic Forecasts. Austrian Fiscal Advisory Council Study. Vienna.

StatCan (2020–2022). Monthly Gross Domestic Product by Industry (Table 36-10-0434-01). Statistics Canada. https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610043401.

Timmer, M.P., Dietzenbacher, E., Los, B., Stehrer, R. and de Vries, G.J. (2015). An Illustrated User Guide to the World Input-Output Database. *Review of International Economics*, 23(3), 575–605.