ETLA

Working Papers | 146

4.9.2026

Growth in Chains: EU Value Chain Productivity and its Slowdown

Tero Kuusi
ETLA Economic Research, Finland
tero.kuusi@etla.fi

Sakari Lähdemäki
ETLA Economic Research, Finland
sakari.lahdemäki@etla.fi

Suggested citation:
Kuusi, Tero & Lähdemäki, Sakari (4.9.2026). "Growth in Chains: EU Value Chain Productivity and its Slowdown". ETLA Working Papers No 146. https://pub.etla.fi/ETLA-Working-Papers-146.pdf

Abstract

We study Europe's post-2008 productivity slowdown using novel total factor productivity (TFP) measures that follow production beyond individual industries in global value chains (GVCs). We collect a newset of stylized facts: (1) the slowdown is broad-based, with productivity growth weakening both in final-producer industries and in upstream industries; (2) the staggered entry into force of bilateral investment treaties with China is followed by higher value-chain TFP and greater business-intangible investment; while (3) the gains associated with technology intangibles are weak. We use these facts to discipline a stylized two-region general-equilibrium model with business capabilities, technology-related knowledge, and cross-border knowledge spillovers. The calibration supports a two-stage interpretation of Europe's experience: (1) sizeable initial integration gains because accumulated business capabilities convert market access into productivity, and (2) later losses when technology improvements are converted into productivity more effectively elsewhere.

Tiivistelmä

Kasvu ketjuissa: EU:n arvoketjujen tuottavuus ja sen hidastuminen

Tutkimme Euroopan vuoden 2008 jälkeistä tuottavuuskasvun hidastumista hyödyntämällä uusia kokonaistuottavuuden mittareita, joiden avulla tuottavuutta voidaan arvioida yksittäisiä toimialoja laajemmin globaaleissa arvoketjuissa. Esitämme joukon uusia empiirisiä havaintoja: (1) Tuottavuuskasvun hidastuminen on laaja-alaista, ja tuottavuuskasvu on heikentynyt sekä lopputuottajatoimialoilla että arvoketjujen tuotantoketjun alkupään toimialoilla; (2) Kiinan kanssa solmittujen kahdenvälisten investointisopimusten voimaantuloa seuraavat arvoketjujen kokonaistuottavuuden kasvu ja liiketoimintakyvykkyyksiin liittyvien aineettomien investointien lisääntyminen; (3) Teknologiaan liittyviin aineettomiin investointeihin liittyvät tuottavuushyödyt ovat vähäisiä. Hyödynnämme näitä havaintoja yleisen tasapainon mallissa, jossa keskeisiä tekijöitä ovat liiketoimintakyvykkyydet, teknologiaan liittyvä tietopääoma sekä rajat ylittävät tiedon leviämisvaikutukset. Analyysi päätyy kaksivaiheiseen tulkintaan Euroopan kehityksestä: (1) Kansainvälisten arvoketjujen integroituessa Euroopan liiketoimintakyvykkyydet mahdollistivat laajemman markkinoille pääsyn muuttumisen tuottavuushyödyiksi, mutta (2) myöhemmin Eurooppa menetti suhteellista asemaansa, koska teknologiaan liittyvät aineettomat investoinnit pystyttiin muualla muuttamaan tuottavuudeksi tehokkaammin.

Ph.D. (Econ.) **Tero Kuusi** is a Research Director at ETLA Economic Research.

D.Sc. (Econ.) **Sakari Lähdemäki** is a Senior Researcher at ETLA Economic Research.

FT **Tero Kuusi** on Elinkeinoelämän tutkimuslaitoksen tutkimusjohtaja.

KTT **Sakari Lähdemäki** on Elinkeinoelämän tutkimuslaitoksen vanhempi tutkija.

Acknowledgements: This work is part of Business Finland's and ETLA's research project "Intangibles as drivers of change and renewal: Firm dynamics underlying the next stage of the knowledge economy" (InChange). ETLA and the authors thank Business Finland for supporting the project. We thank Gaaitzen de Vries, Roberto Samaniego, Adele Whelan, and Qizhong Yang for their comments, and attendees at ASSA Annual Meeting 2026, the 8th World KLEMS, CUFE-OEP, Euroframe 2026, FEBST 2025, TSER, the Finnish Economic Association's 45th Annual Meeting, and the 40th Summer Seminar of Finnish Economists for their valuable feedback.

Artificial intelligence has been used to support human work in the preparation of this working paper in accordance with Etla's ethical guidelines (version 14.1.2026, see https://www.etla.fi/en/ai-ethics).

Kiitokset: Tämä tutkimus on saanut rahoitusta Business Finlandin ja Etlan InChange -tutkimushankkeesta. Etla ja kirjoittajat kiittävät Business Finlandia rahoituksesta. Kiitämme lisäksi Gaaitzen de Vriesiä, Roberto Samaniegoa, Adele Whelania ja Qizhong Yangia kommenteista sekä ASSA Annual Meeting 2026 -konferenssin, 8th World KLEMS -konferenssin, CUFE-OEP:n, Euroframe 2026:n, FEBST 2025:n, TSER:n, Taloustieteellisen Yhdistyksen 45. vuosikokouksen ja 40. Kansantaloustieteen kesäseminaarin osallistujia arvokkaasta palautteesta.

Tämän työpaperin työstämisessä on hyödynnetty tekoälyä ihmistyön tukena Etlan eettisen ohjeiston mukaisesti (versio 14.1.2026, ks. https://www.etla.fi/ai-etiikka).

Keywords: Global value chains, Intangibles, Productivity, Globalization

Asiasanat: Globaalit arvoketjut, Aineettomat investoinnit, Tuottavuus, Globalisaatio

JEL: F60, O47, L16, O14, E22

1 Introduction

Despite rapid advances in digital technologies, continued diffusion of new organizational practices, and historically deep integration into global markets, total factor productivity (TFP) growth across Europe has slowed markedly after 2008. Policy discussions increasingly stress that Europe's innovation and productivity problems are intertwined with limited scale and incomplete integration (Draghi, 2024). Moreover, research points to complementary intangible investments and organizational reconfiguration that technologies require before productivity gains appear (Brynjolfsson et al., 2021). However, the productivity impact of business innovation-related intangibles has become more complex over time and increasingly hinges on complementarities with other forms of capital and on diffusion conditions (van Ark et al., 2024). These perspectives help rationalize the European problems, but they are hard to quantify: Through which parts of GVCs do these forces operate and how they are interrelated?

In this paper, we make empirical and quantitative theoretical contributions in answering these questions. A core difficulty is measurement. While globalization fragments production across borders and industries, conventional productivity measures are not designed to track how productivity evolves within value chains or how upstream and downstream performance co-move with intangible accumulation and trade integration.[1] Thus, we develop and apply a value-chain productivity metric that is explicitly designed to track productivity dynamics at the level where globalization operates: the value chain rather than the standalone industry.

Our empirical analysis builds a unified set of facts about European productivity in chains, involving impacts of globalization shocks and the total-factor productivity elasticities of technology investments. We find that Europe's slowdown looks different when measured at the value-chain level. We document that Europe's post-crisis productivity slowdown reflects weaker VC TFP growth both in final producer industries and in the upstream segments of their value chains.

[1] The traditional measurements are based on single industries and require price deflators for intermediate inputs that are notoriously hard to measure in internationally fragmented production (e.g., due to quality change, transfer pricing, and the valuation of intangible-rich inputs) (Houseman and Mandel, 2015; Timmer, 2017; Timmer and Ye, 2020).

We then use this VC TFP panel in three complementary analyses that together build a causal account of the dynamics.

We first analyze evidence on the productivity impacts of globalization shocks, with the focus on the early-2000s investment liberalization between China and EU15 countries. We implement an event-study design around the entry into force of bilateral investment treaties (BITs) using treatment-timing robust methods (Callaway and Sant'Anna, 2021). The results indicate that BITs are followed by a statistically and economically significant increase in VC TFP in the market sector, with effects visible both in final producer industries and in upstream segments of the value chain. We also find that business intangibles in the final producer industry increase after the treaty comes into force, while technology-related intangibles show a weaker and noisier response. To assess whether these aggregate effects mask heterogeneity, we split industries at the median of their pre-treaty Chinese value-added share, forming high- and low-China-exposure subsamples. We then estimate the CS-DiD event study separately in each group. Both partitions display positive long-run BIT effects on TFP and intangibles, and the high-versus-low differences are not statistically distinguishable from zero, indicating that the aggregate response is broad-based rather than concentrated in the most China-integrated industries.

To assess what the observed changes in intangible investment imply for VC TFP more systematically, and to provide elasticities that can discipline further modeling exercises, we estimate dynamic panel models relating VC TFP growth to log intangible investments, following van Ark et al. (2024), Corrado et al. (2022), and Corrado et al. (2017). The event study already points to a strong business-intangible–TFP link, while we further analyse this relationship through panel elasticity estimates that condition on additional controls and exploit within-chain variation over time. Our evidence points to a robust pre-2008 association between business intangibles and VC TFP growth. The estimated association for business intangibles is substantially attenuated in the post-2008 period. These relationships emerge only with the VC TFP measure: repeating the same specification with conventional EU KLEMS TFP as the outcome leaves the business-intangible coefficient statistically indistinguishable from zero.

Finally, we use these facts to discipline a stylized two-region general-equilibrium model of trade and intangibles, which serves as an interpretive framework rather than as a second source of evidence. It features Armington-CES trade in manufactured goods, business capabilities that are accumulated jointly with production, technology-related knowledge, and cross-border knowledge spillovers in both types of intangible.

The exercise separates two mechanisms that the empirical work cannot separate on its own. First, lower integration frictions expand Europe's market access and, because accumulated business capabilities convert that access into productivity, raise European value-chain productivity while simultaneously reallocating manufacturing toward China. Second, matching the empirical moments requires the elasticity of TFP with respect to technology-related knowledge to be substantially larger in China than in Europe, so the same global technology improvements shift economic advantage toward China over time. Taken together, the paper's lesson is that the durability of Europe's gains from globalization depends not only on access to global production networks but on the interaction between organizational capability and technology absorption: Europe can hold an initial advantage in business capabilities and still experience a relative productivity slowdown if technology adaptation is stronger elsewhere.

The remainder of the paper is organized as follows. Section 2 reviews related literature. Section 3 describes the data and the construction of value-chain variables. Section 4 documents the rise and fall of EU value-chain productivity and decomposes VC TFP into final-producer versus upstream components. Section 5 uses an event-study design to study globalization shocks through bilateral investment treaties with China, including heterogeneity analysis by industry exposure, and estimates reduced-form elasticities between intangible investment and VC TFP. Section 6 presents a stylized two-region model of trade and intangibles, motivated by the GVC evidence, that organizes the empirical patterns and draws out their implications for the post-2008 divergence between China and Europe. Section 7 concludes. The formal model environment, additional heterogeneity analysis, value-added composition, and robustness checks are reported in Appendix A.

2 Related literature

This paper links to several strands of literature. First, it relates to work on European competitiveness. The Draghi report frames the European productivity slowdown as a problem of scale, diffusion, and fragmentation. It argues that Europe's productivity challenge is tightly linked to the ability to build scale in key activities and to diffuse innovations widely across a still-fragmented single market (Draghi, 2024). A related work highlights the role of scale and policy in shaping comparative advantage and the returns to capability accumulation. Juhász et al. (2024) review new empirical evidence on industrial policy mechanisms—learning, coordination, and scale-related externalities. Cuñat and Zymek (2025) study how persistent global imbalances can interact with scale economies and induce durable reallocations in tradable production. These ideas map naturally to our emphasis on production scale as a mediator of production effectiveness.

Second, our focus on intangibles[2] is consistent with the idea that productivity growth in modern economies depends not only on invention, but on broad-based adoption, organizational upgrading, and the ability to operate at scale. Brynjolfsson et al. (2021) argues that general-purpose technologies require complementary intangible investments and reorganization, which affect (measured) productivity even when underlying capabilities are improving. Guzman et al. (2024) emphasize that regional innovation ecosystems and the policy environment surrounding them shape how innovative capability is built.

There is evidence showing that intangible capital has contributed positively to productivity growth (Corrado et al., 2022; Corrado et al., 2017; and van Ark et al., 2024). In developed countries, there tends to be a shift towards more non-routine and more interactive tasks, and the use of highly educated workers (Becker et al., 2013; Reijnders and de Vries, 2018), which benefits industries that are intangible intensive (Jona-Lasinio and Meliciani, 2019). Corresponding to our time period, Deng and Liu (2024) show that Italian firms reallocated from intangible assets to

[2]In terms of measurement of intangible investment and productivity accounting, we also draw on the EU KLEMS & INTANProd framework and documentation, which provide a harmonized classification of intangible assets and discuss key measurement challenges (Bontadini et al., 2023). By studying VC TFP, our work complements the research efforts of van Ark et al. (2024) that uses the database to study the contribution of intangible capital deepening to labor productivity growth and TFP spillovers from intangible capital.

tangible assets during the 2011–2012 sovereign-debt crisis, affecting adversely both output and TFP.

More directly relevant for the translation of capability into measured productivity is the technology-diffusion and absorptive-capacity literature: Keller (2004) stresses that international diffusion is neither automatic nor costless and depends on domestic technology investment, Griffith et al. (2004) document the "two faces" of R&D as both innovation and absorptive capacity, and Bloom et al. (2012) show that management practices condition the productivity payoff to information technology in European establishments. In the European context, van Ark et al. (2024) argues that the productivity payoff to business innovation-related intangibles has become more nuanced and more dependent on complementarities with other assets and with diffusion conditions. Relative to this strand, our contribution is to distinguish business from technology intangibles and to show that their association with value-chain TFP changes sharply after 2008—a pattern that these insights help interpret.

Several articles have studied the motives for offshoring arising from the interaction of several complex factors (Baldwin and Robert-Nicoud, 2014). The theoretical architecture of multi-stage global value chains has been analyzed by Lee and Yi (2018), who show that multi-stage global production can magnify the aggregate effects of lower trade costs, because integration affects several linked production stages at once. Our framework abstracts from explicit stage-by-stage specialization and worker sorting. Instead it focuses on a complementary mechanism: how expanded market access interacts with business capabilities and cross-border technology diffusion, allowing the productivity response to integration to differ across regions. Relative to this theoretical strand, then, our model deliberately collapses multi-stage production in order to concentrate on capability accumulation and technology diffusion, and it should not be read as a simplified replica of a full sequential-production model. Our work also contributes to the previous literature by showing more evidence that growing international trade and better organization of international production into global value chains led to productivity gains in the 2000s (Constantinescu et al., 2019; Goldin et al., 2024).

Zeng and Lu (2016) documents the determinants of BIT formation with China. Fan et al. (2018) show that China's tariff reductions facilitated quality upgrades, raising the efficiency of Chinese firms' use of intermediate inputs. Chor et al. (2021) document a notable rise in the upstreamness of imports, a stable positioning of exports, and a rapid expansion of production stages conducted in China from 1992 to 2014—both at the aggregate level and within firms over time.

Our event study analysis relates with studies of the broader impacts of China (and other countries) on local market dynamics through international trade, see for example, Autor et al. (2013) and Autor et al. (2016). Relative to that literature, we use treaty timing to study European value-chain productivity rather than local labour-market exposure or Chinese firm upgrading. Identifying the causal effects of China integration on domestic industries requires a source of plausibly exogenous variation in exposure to value chains. This work is also related to the literature on measuring shocks in value chains (see, e.g., Baqaee and Farhi, 2019; Acemoglu and Tahbaz-Salehi, 2020; Elliott et al., 2022), and their empirical propagation (Bonadio et al., 2021; Boehm et al., 2019; Carvalho et al., 2021).

In terms of measurement of productivity, our efforts are primarily driven by the challenges in accurately measuring the prices of intermediate inputs, which stem from transfer pricing practices in multinational enterprises, the complexities in valuing intangible flows, and the inadequate statistical systems for tracking intermediate prices amid quality improvements (Houseman and Mandel, 2015; Timmer, 2017; Goldin et al., 2024). We use a decomposition of the global value-added contents of the outputs and contributions of industries and of the other sectors in the upstream value chain (Leontief, 1936; Wolff, 1994; Timmer, 2017; Timmer and Ye, 2020, Kuusi et al., 2022). The value chain approach makes visible both the substantial role of upstream industries to which industries have backward linkages as well as technology and knowledge investments as a source of productivity growth in the entire value chain (for a review, see Carvalho et al. (2021)). Relative to this literature, our contribution is to use a value-chain TFP measure not only for accounting but as the outcome variable in a causal integration design and as an input to a quantitative interpretation of the European slowdown.

3 Data

To form our VC TFP panel, that covers the years 1995-2017 and includes 675 VC TFP series for the EU15 industries, we combine data from OECD, World KLEMS, The Conference Board and Penn World Tables 10.1. Two key datasets are: OECD Inter-Country Input-Output Tables (ICIO) and World KLEMS data (KLEMS). ICIO input-output tables cover 76 individual countries (38 OECD countries and 38 non-OECD economies) and an aggregate for the rest of the world (ROW). For each country it provides input-output data for 45 industries. KLEMS provides input factor data (labor and capital inputs and shares) for 30 countries and basically the same industries.[3] Additionally, we gather input factor data for 7 other countries from OECD STAN database, which has the same industry classification as ICIO. For those countries for which we do not have industry level factor inputs data we use country-level data from Penn world tables 10.1 (PWT10.1).[4] Finally, we calculate aggregate input factor and weighted average input share series for ROW using PWT10.1 that provides input factor data in levels.

For some countries we impute missing input factor and input share values for the years 1995-2008.[5] Furthermore, following Kuusi et al. (2022), if we lack data for certain industries in a certain country we use the "next" aggregate level data as a proxy of this industry.[6] This is reasonable because we deal with growth rates and input shares. Undeniably, this adds some inaccuracy to our

[3]We use the ICIO industry classification in our dataset and adjust the industry division in the KLEMS data accordingly. Note that for China the industry classification is somewhat different from EU KLEMS data. In this case the most closest industry (or the nearest higher level aggregate) is used as a proxy for the ICIO industry.

[4]PWT10.1 lacks labor and capital share data for BGD, KHM, MMR, PAK and VNM. For these countries this data is from The Conference Board (TCB). Additionally, both PWT10.1 and TCB lack data for BRN, for this country we use the input factors and shares calculated for ROW.

[5]For BEL (years 1995-1998) and BGR, HRV, CYP, EST, GRC, HUN, IRL, LVA, LTU, LUX, MLT, POL, PRT, ROU, SVN (years 1995-2007) the KLEMS labour services volume indices are continued using the growth rates of persons employed. For USA the missing values for year 1995 of KLEMS capital service volume index are set to be the same as in 1996. For CAN capital and labor shares are set to be the same as in year 2007 (years 1995-2006) and for the years 1995-1996 in some industries the missing data for total employment are replaced by the value in year 1998. For CHL the missing values for year 1995 of OECD Stan net capital stock are set to be the same as in 1996 and capital and labor shares are set to be the same as in year 2007 (years 1995-2006). Note also that in some rare cases, for some specific industry and year, there are negative value added values in the data. In these cases, we replace the negative value with a near-zero positive value to enable the use of logarithms.

[6]For example, if data for petroleum and coal products (19) is missing for certain country, we use data for chemical, rubber, plastics, fuel products and other non-metallic mineral products (19-23) and if this data is also missing we then use data for whole manufacturing (c). Ultimately, we use data for total value added if no other closely related aggregate data is available.

measures, but the next-level data can be considered a good proxy because the growth rates of a certain industry are likely highly correlated with the next aggregate-level data. In Appendix A, we carefully scrutinize the differences between traditional TFP and VC TFP.

Out of 76 countries we have meaningful input data at industry-level for 37 countries.[7] For countries for which no industry-level data are available we use country-level input data. The 37 countries for which we have input data at industry-level include the most developed and largest economies of the world (also new KLEMS data for China industries). Therefore, arguably, for the EU15 countries we study, the input factor and share data is available at industry level for most parts of their global value chains. Therefore, regardless of the data shortages, for the EU15 the measure of global value chain TFP can be considered a good approximation of the "true" global value chain TFP.

Other variables, such as traditional TFP and output are from KLEMS. The intangible variables are also from KLEMS and the division of total intangibles to technology related and business innovation related intangibles is calculated following van Ark et al. (2024) and Corrado et al. (2022).[8] Data on Bilateral Investment Treaties (BIT) are from UNCTAD.

4 European VC productivity growth

In this section we calculate industry-level TFP that accounts for all input factors used in the industry's value chain (VC TFP). The calculation of VC TFP relies on two key methodological achievements in the literature. Firstly, we calculate value added contributions of the global value chain using the hypothetical extraction method by Los et al. (2016). Accordingly, the difference between the actual global value added distribution and a hypothetical distribution, where the production of a certain industry is set to zero, gives the total value added in the value chain of this industry.

[7] Note that while for most countries the data covers the different industries rather comprehensively there are some exception. For example, for CHL we only have industrial data on a high level of aggregation.

[8] Technology related intangibles are formed by summing the real values of computer software and databases, R&D, entertainment, artistic and literary originals and mineral exploration from KLEMS intangibles data. Accordingly, business innovation related intangibles include new product development costs in the financial industry, industrial design, brand, organizational capital and training.

Secondly, following Timmer (2017) and Kuusi et al. (2022), using these industry's value added contribution to the global value chain and country-industry specific input factors we calculate input factor requirements of a certain industry. Finally, the input factor requirements and country-industry specific input shares allow us to derive VC TFP for a certain industry from a traditional production function.

We next describe the derivation of VC TFP in more detail, for more see Los et al. (2016), Timmer (2017) and Kuusi et al. (2022). The hypothetical extraction method is as follows.

$$VA = v(I - A)^{-1}Y * i \tag{1}$$

where VA is a value chain matrix that contains industry- and country specific value added contributions, v is a row vector of the ratios of value added to gross output, i is a column vector of ones implying that it sums all elements of the final demand matrix Y. $(I - A)^{-1}$ is the Leontief inverse where I is a identity matrix and A is a matrix that contains the use of intermediate products.

The hypothetical world where certain country-industry's production is excluded can be derived by extracting this industry' production of final goods and intermediate products to other industries in own and other countries and calculate VA again. The hypothetical world is denoted VA^*. Then,

$$\Delta VA = VA - VA^* \tag{2}$$

ΔVA contains the value added contributions of different industries of different countries to the specific industry in question. Further on we denote these value added contributions $VAVC$ and each contributions share as VAs. Value chain TFP of country-industry J can be obtained as follows (for brevity time index t is dropped):

$$\Delta \log(Q_J) = \sum_{j=1}^{3465} \overline{VAs}_{Jj}(\overline{\alpha}_L^j * \Delta \log(\tilde{L}_j) + \overline{\alpha}_K^j * \Delta \log(\tilde{K}_j)) + \Delta \log(TFP_J) \tag{3}$$

where

$$\tilde{L}_j = \frac{L_j VAVC_{Jj}}{VA_j}, \qquad \tilde{K}_j = \frac{K_j VAVC_{Jj}}{VA_j}. \tag{4}$$

In equation (3) Q is output in real terms, $\tilde{L}$ is labor and $\tilde{K}$ is capital requirements and TFP is total factor productivity. $\overline{VAs_{Jj}}$ is the Törnqvist share of value added contribution's yearly share of country-industry j from the country-industry's J value chain. In equations (4) L_j is labor input and K_j is capital input in real terms of country-industry j. VA_j is total value added of country-industry j and $VAVC_{Jj}$ is the value added contribution of country-industry j from the country-industry's J value chain. $\overline{\alpha}_{L,K}$ are the Törnqvist shares of the yearly labor and capital shares of country-industry j. We can solve TFP from equation (3) because we have data on everything else.

4.1 Value chain growth dynamics

We begin by documenting value chain productivity dynamics in EU15 industries over 1995–2017 and how these dynamics changed after the 2008 crisis.[9] Table 1 reports output growth and its decomposition into input contributions and VC TFP growth for broad sectors. Two aggregate facts motivate the rest of the paper. First, before 2008, EU15 value chains exhibited positive VC TFP growth, with manufacturing in particular showing sizeable VC TFP contribution. Second, after 2008, output growth stagnated and VC TFP growth weakened broadly, consistent with a Europe-wide slowdown in productivity growth within global value chains.

VC TFP aggregates productivity across the entire chain, so it is useful to ask whether the slowdown originates primarily in final producer industries or in the upstream part of the chain. Following Timmer (2017), we use EU KLEMS TFP growth estimates to decompose VC TFP into the final producer industry's TFP growth and the TFP growth originating in the rest of the value chain. The two components are defined as follows. The final-producer component is the standard

[9] Appendix A.1 validates the VC TFP measure by comparing annual VC TFP growth g^{VC} to conventional EU KLEMS TFP growth g^{KLEMS} via the signed wedge $\varepsilon = g^{VC} - g^{KLEMS}$ and its square, ε^2 (Appendix Tables 9 and 10). Three patterns stand out. First, the wedge is systematically larger in value chains with a larger market-services component, consistent with the idea that conventional industry-based TFP measures are particularly restrictive in service-intensive and networked value chains. Second, neither the level nor the volatility of the wedge rises systematically with the foreign share of the value chain, so the use of foreign upstream information does not mechanically create extra volatility or systematic bias. Third, the results are consistent with earlier exercises in the literature.

Table 1: Output growth and its decomposition into TFP and input growth contribution within value chains for different sectors (rows)

	1995-2007 Output growth	Input contr.	VC TFP	2008-2017 Output growth	Input contr.	VC TFP
Manufacturing	2.5	1.3	1.2	-0.7	-0.8	0.1
Market serv	4.6	3.8	0.8	1.1	1.6	-0.5
Other serv	1.8	2.2	-0.4	0.9	1.5	-0.6
Utilities + cons	2.8	2.5	0.3	0.0	0.8	-0.8
Average	**3.2**	**2.4**	**0.8**	**0.1**	**0.4**	**-0.2**

Notes: Entries are average annual growth rates (percentage points). Output growth is decomposed into the contribution of aggregate primary inputs used across the full value chain and value-chain TFP (VC TFP). EU15; industries grouped into broad sectors.

value-added-based TFP growth measure for the final industry itself, taken from the conventional EU KLEMS growth-accounting framework. We then attribute the remaining value-chain productivity contribution to the upstream country–industry components supplying the chain; operationally this “other VC TFP” component is computed as the residual, so that by construction the two components sum to the VC TFP measure reported above. This makes the decomposition directly comparable to conventional growth accounting while retaining the value-chain perspective. Table 2 shows that the post-2008 slowdown is visible in both components. In manufacturing and market services, the upstream part of the chain contributes little to TFP growth pre-2008 and turns negative after 2008, while final producer TFP growth also weakens.

Table 2: TFP decomposition to final industry TFP and the rest of the value chain TFP by sector (rows)

	1995-2007 Final TFP	Other VC TFP	2008-2017 Final TFP	Other VC TFP
Manufacturing	1.0	0.2	0.3	-0.1
Market serv	0.8	0.1	0.0	-0.5
Other serv	-0.6	0.2	-0.4	-0.2
Utilities + cons	0.0	0.3	-0.6	-0.2
Average	**0.7**	**0.2**	**0.0**	**-0.3**

Notes: Value-chain TFP is decomposed into the TFP of the final-producer industry and the residual TFP from the upstream part of the value chain supplying intermediate inputs. Entries are average annual growth rates (percentage points). EU15.

The aggregate decomposition masks meaningful differences in where the slowdown appears. Country-level decompositions (Appendix Tables 12 and 13) reveal substantial cross-country variation: some countries maintained positive VC TFP growth after 2008, often through input contraction, while others experienced broad TFP weakness across sectors. The decomposition of VC TFP into final-producer versus upstream contributions (Appendix Table 14) shows that the post-2008 upstream deterioration is especially pronounced in market services, that is, in precisely the segments where conventional industry-level accounting is least informative about chain-wide productivity. Input composition (Appendix Tables 15 and 16) documents a broad contraction in both domestic and foreign input growth after 2008, particularly in manufacturing.

Value-added composition analysis (Appendix Tables 17 and 19) shows that market services account for about a quarter of value added in European value chains and that their role has grown over time, particularly through the foreign upstream component. Taken together, these patterns indicate a broad shift in which European value chains became more service- and import-intensive at the same time as their measured productivity growth deteriorated. This finding motivates us to focus on mechanisms that operate at the level of chain organization and intangible capability rather than within single industries.

5 Analyzing the value-chain productivity

This section presents the paper's main empirical analysis. We first ask whether the entry into force of bilateral investment treaties (BITs) with China is followed by changes in value-chain productivity and intangible investment, and whether the response varies with an industry's pre-treaty exposure to China. We then estimate panel elasticities that condition on additional controls and exploit within-chain variation, which supply the elasticity moments used in Section 6.

5.1 Empirical strategy

Our first empirical strategy uses policy-driven liberalization events that can shift the organization and scale of cross-border production. We focus on China's bilateral investment treaties with EU countries signed in the period 2000–2014 and treat the years when the treaties enter into force as the event years.[10] These treaties expanded the legal security of bilateral investment relationships and included broader two-way provisions, allowing investor-to-state dispute settlement.

Treatment is assigned at the country level and the outcomes are measured at the country–industry level. Given staggered timing and likely treatment-effect heterogeneity, we use the event-study estimator of Callaway and Sant'Anna (2021). As a baseline, the control group consists of countries where only first-generation BITs with China are in force. The treated countries are the Netherlands (2004), Sweden (2004), Germany (2005), Finland (2006), Portugal (2008), Spain (2008), Belgium (2009), Luxembourg (2009), and France (2010), while the baseline control group includes Austria, Denmark, Greece, Ireland, Italy, and the United Kingdom. We focus on an event window from seven years before to seven years after the event so that all treated countries contribute to each coefficient estimate; we therefore refer to the t+7 ATT as the long-run effect throughout the analysis.

Inference is based on resampling at the level of the country–industry panel unit; this is the baseline inference throughout the paper, and all figure and table notes refer to it. Because BIT timing is assigned at the country level, the coarser country-level clustering that matches the level of treatment assignment is the natural alternative, and we report it explicitly as robustness in Appendix A.3 together with additional robustness exercises, including treated-only countries, alternative windows, and alternative estimators. We additionally examine heterogeneity in treatment effects by partitioning the sample at the industry-level median of each industry's pre-determined exposure to Chinese value chains and estimating the CS-DiD event study separately in the high- and low-exposure partitions; the complete methodology is described in Appendix A.3.

[10] Another possibility would be to use the signing year as the event year. Our pre-trend analysis, however, suggests that entry into force is the more suitable timing, and Zeng and Lu (2016) document substantial variation in the content and legal strength of China's BIT provisions, which is consistent with treaty effects depending on the point at which commitments become binding.

Our second empirical strategy studies the relationship between intangible investment and VC TFP growth. Following van Ark et al. (2024), Corrado et al. (2022), and Corrado et al. (2017), we estimate panel models that relate productivity growth to the current and lagged *log intangible stocks* in business and technology-related intangibles, tangible input growth, and the foreign market-services share. We also include an interaction between business and technology intangibles. We use a standard dynamic specification: the contemporaneous and lagged stock coefficients together capture the growth rate of the stocks.[11]

Let g^{VC}_{ict} denote annual VC TFP growth in country–industry–year cell (i, c, t), then our model is

$$g^{VC}_{ict} = \beta_A b_{ict} + \beta_{A,-1} b_{ic,t-1} + \beta_Z z_{ict} + \beta_{Z,-1} z_{ic,t-1} \tag{5}$$

$$+ \beta_{AZ}(b_{ict} z_{ict}) + \beta_{AZ,-1}(b_{ic,t-1} z_{ic,t-1}) + \beta_Q q_{ict} + \beta_{Q,-1} q_{ic,t-1} + \mu_{ic} + \tau_t + u_{ict}, \tag{6}$$

where $b_{ict} = \ln(\text{bus. intangible stock}_{ict})$ and $z_{ict} = \ln(\text{tech. intangible stock}_{ict})$ are log intangible capital stocks, q collects tangible-input growth and the foreign market-services share, μ_{ic} is a country–industry fixed effect, and τ_t is a year effect.

We estimate this specification by fixed effects OLS (FE) with robust standard errors clustered at the country–industry panel identifier as the baseline. The FE columns are the primary evidence. The pre-2008 subperiod provides the cleanest identification, as the post-2008 pooled specifications are affected by the crisis-induced regime change documented below. We report a system-GMM specification alongside the FE results in Table 4 (column (3)). Our static GMM—which excludes the lagged dependent variable from the regression—uses a collapsed instrument matrix (lag window [2,5], 31 instruments) and passes the AR(2) test. Additional robustness specifications are reported in Appendix A.4 and confirm the main findings.

[11] Using log stocks with a one-period lag, the contemporaneous and lagged coefficients are approximately equal and opposite in the level-FE specification—a pattern we verify in the data—so the stock-level regression effectively recovers a relationship between investment flows and TFP growth, proportional to net investment.

5.2 Results

Globalization shocks and value-chain productivity. We first consider the response of productivity to BIT-induced integration with China. We concentrate on the market sector (manufacturing and market services), which is the part of the economy most likely to respond through changes in value-chain organization. Figure 1 reports dynamic effects on traditional TFP and VC TFP. Both measures increase after the treaties enter into force, but the increase is more pronounced and more clearly sustained for VC TFP. The pre-treatment coefficients fluctuate close to zero, which supports the identifying assumption. The same globalization shock is more clearly visible when productivity is measured at the value-chain level than industry by industry.

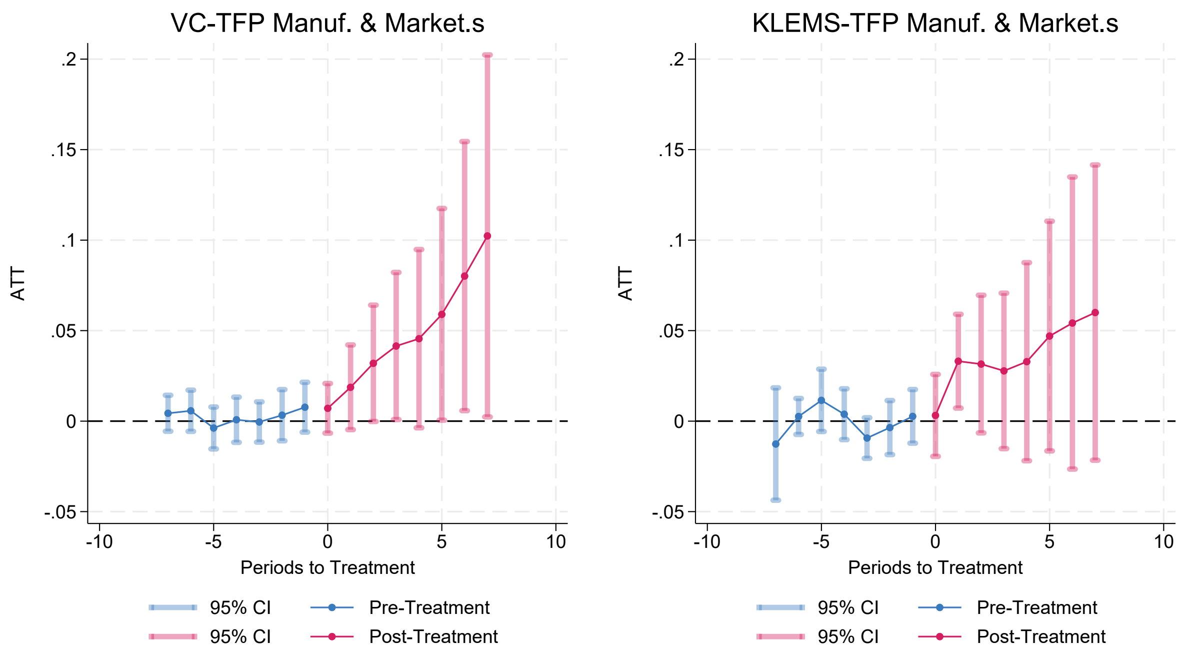

Figure 1: Development of TFP after BITs agreement enters into force

Note: The dots are event coefficients and the bars their bootstrapped 95 % confidence intervals clustered at unit level. The pre-treatment coefficients are blue and the post-treatment coefficients red.

Figure 2 then decomposes VC TFP into productivity growth in the final producer industry and in the rest of the value chain. The positive post-treaty response is visible in both components, although it is more pronounced in the final producer industry. This suggests that the benefits of integration do not operate solely through the final industry itself, but also through reorganization and productivity changes in upstream segments of the chain.

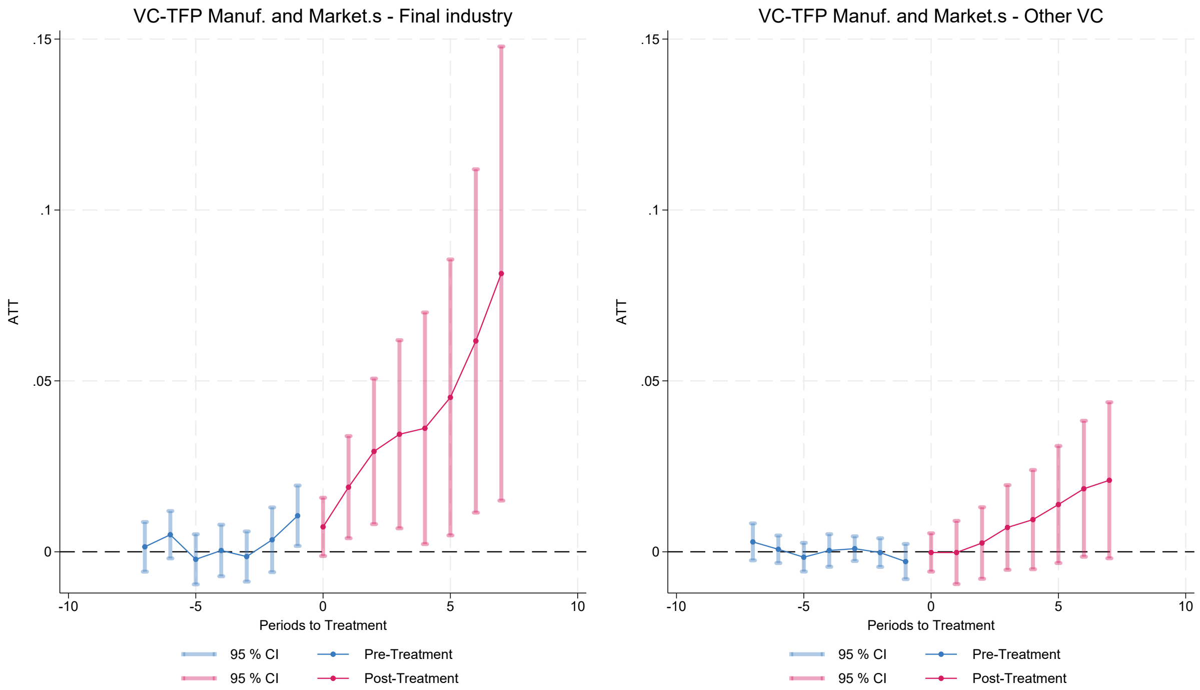

Figure 2: Decomposition of VC TFP after BITs agreement enters into force: final producer vs. upstream chain

Note: The dots are event coefficients and the bars their bootstrapped 95 % confidence intervals clustered at unit level. The pre-treatment coefficients are blue and the post-treatment coefficients red.

The mechanism is further clarified in Figure 3. Output and value added rise after the treaty enters into force, while the increase in total factor requirements along the value chain is smaller. This is consistent with the idea that BITs improved the organization and scale of cross-border production, allowing output to expand faster than value-chain inputs. The same figure also shows only limited movement in the domestic value-added share, suggesting that the main effect operates through productivity-enhancing reorganization rather than a simple substitution between domestic and foreign stages.

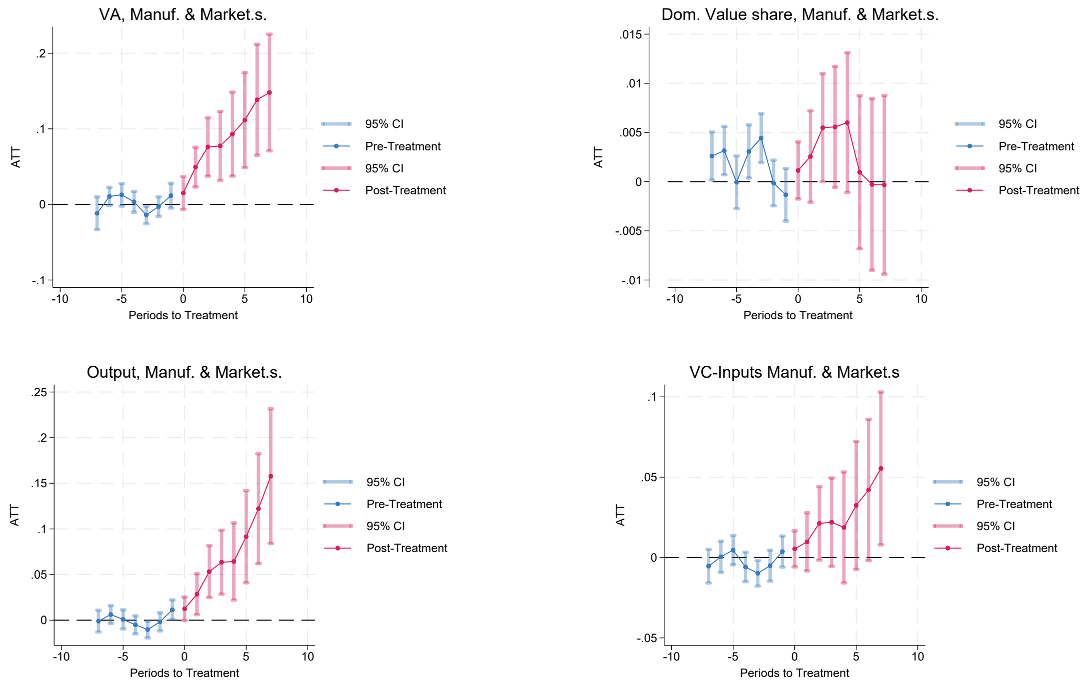

Figure 3: Development of value added, output, inputs, and domestic value share after BITs agreement enters into force

Note: The dots are event coefficients and the bars their bootstrapped 95 % confidence intervals clustered at unit level. The pre-treatment coefficients are blue and the post-treatment coefficients red.

A natural candidate mechanism is intangible investment in the final producer industry. Figure 4 shows that business-related intangible investment flows increase after BITs enter into force, while the corresponding response of technology-related intangible flows is visibly weaker and noisier in the available sample.

The results indicate that the strongest reduced-form response to globalization is not a pure technology-adoption effect, but rather an increase in business-oriented investment that is plausibly linked to reorganization, coordination, and the scaling up of activity in the chain.

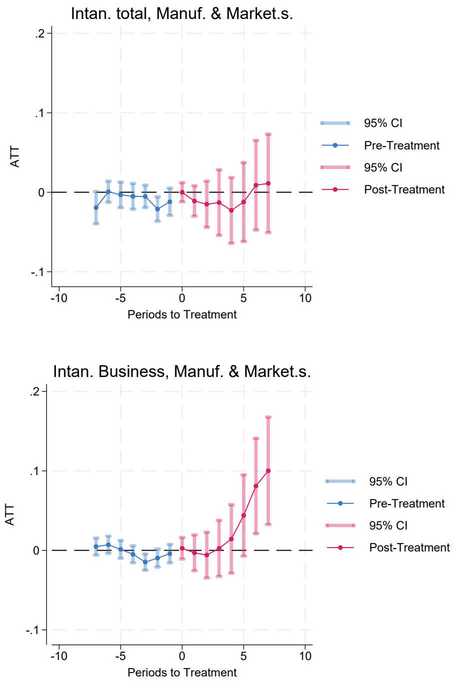

Figure 4: Development of final-industry intangible investment flows after BITs enter into force

Note: The dots are CS-DiD event coefficients (Callaway–Sant'Anna doubly-robust estimator; Callaway and Sant'Anna, 2021; Sant'Anna and Zhao, 2020) and the bars their bootstrapped 95 % confidence intervals (999 Mammen wild bootstrap replications, resampled at the country–industry panel-unit level). Treatment timing is assigned at the country level, so identification rests on 15 countries; country-level clustering is reported as robustness in Appendix A.3. The outcome is log annual intangible investment expenditure (flow). The pre-treatment coefficients are blue and the post-treatment coefficients red. Sample size: 8,818 obs for business intangibles, 6,158 for technology intangibles.

Table 3 presents the full ATT sequence at each horizon $k = -7$ to $k = +7$ for cumulative VC TFP growth, log business intangible investment, and log technology intangible investment. The pre-BIT coefficients for VC TFP are small and statistically insignificant throughout, supporting parallel trends for that outcome. For business intangibles, the pre-treatment estimates are also modest but show a mild negative tendency in the two years immediately before BIT entry ($k = -3$: -0.015; $k = -2$: -0.014), possibly reflecting anticipatory reallocation. The post-BIT trajectory shows a smoothly rising profile for both VC TFP and business intangibles, reaching $\hat{\tau}_{+7}^{TFP} = 0.102$ (SE

= 0.060, p = 0.087) and $\hat{\tau}^{bus}_{+7}$ = 0.105 (SE = 0.032) at the long-run (t+7) horizon—the maximum horizon at which all treated countries contribute to the ATT estimate. Technology intangibles present a more uncertain picture: the pre-treatment estimates are noisy with statistically significant deviations at $k = -5$ (0.041, SE = 0.019) and $k = -2$ (−0.046, SE = 0.014), raising parallel-trends concerns for this outcome. The post-treatment profile is non-monotone and the seven-year estimate $\hat{\tau}^{tech}_{+7}$ = 0.030 (SE = 0.043, p = 0.486) is statistically insignificant, consistent with Figure 4.

Table 21 in Appendix A.2 formalizes this assessment via a joint Wald test of all six pre-treatment ATTs ($k = -7$ through $k = -2$). For VC TFP the joint null is not rejected ($\chi^2(6) = 8.22$, $p = 0.22$), and for business intangibles the test likewise fails to reject at conventional levels ($\chi^2(6) = 12.07$, $p = 0.06$), supporting the parallel-trends assumption for the two outcomes that enter the model calibration. For technology intangibles the joint test rejects, which is consistent with the individual pre-period deviations noted above and with the insignificant post-treatment ATT for that outcome. Since technology intangibles are not used as a calibration moment, this pre-trend violation does not compromise the structural exercise.

The main limitation of this design is that, although the wild bootstrap resamples at the country–industry panel-unit level, treatment timing varies across only 15 countries (9 treated and 6 control). Inference therefore remains approximate at the level at which the policy actually operates, and splitting the sample by industry does not create additional country-level policy experiments: industries within a country share the same BIT timing and are exposed to common country shocks. As an additional robustness check, we also restrict the sample to countries that eventually signed a BIT with China. The resulting estimates remain positive and of similar magnitude, which suggests that the baseline findings are not driven simply by persistent differences between BIT-signing and non-signing countries. At the same time, as discussed in the Appendix, this treated-only exercise should not be interpreted as eliminating all endogeneity concerns in treaty timing.

Differential effects by industry exposure. The country-level event study captures the *average* treatment effect over all industries. To assess whether this aggregate masks systematic hetero-

Table 3: Average treatment effects at each horizon: VC TFP and intangible investment

	VC TFP (cumul.)		Bus. intangibles (log flow)		Tech. intangibles (log flow)	
Horizon	ATT	SE	ATT	SE	ATT	SE
Pre-treatment periods						
$k = -7$	0.004	(0.004)	−0.001	(0.006)	−0.049	(0.031)
$k = -6$	0.006	(0.006)	0.004	(0.005)	0.018	(0.013)
$k = -5$	−0.004	(0.007)	0.004	(0.005)	0.041	(0.019)
$k = -4$	0.001	(0.005)	−0.004	(0.005)	−0.025	(0.019)
$k = -3$	−0.000	(0.005)	−0.014	(0.005)	0.005	(0.020)
$k = -2$	0.003	(0.006)	−0.014	(0.005)	−0.046	(0.014)
Post-treatment periods						
$k = +0$	0.007	(0.007)	0.014	(0.007)	0.008	(0.014)
$k = +1$	0.019	(0.014)	0.020	(0.012)	0.028	(0.015)
$k = +2$	0.032	(0.021)	0.020	(0.015)	0.057	(0.026)
$k = +3$	0.042	(0.028)	0.034	(0.018)	0.073	(0.036)
$k = +4$	0.046	(0.034)	0.037	(0.022)	0.021	(0.035)
$k = +5$	0.059	(0.040)	0.060	(0.025)	0.021	(0.037)
$k = +6$	0.080	(0.048)	0.084	(0.028)	0.020	(0.039)
$k = +7$	0.102	(0.060)	0.105	(0.032)	0.030	(0.043)

Notes: Callaway–Sant'Anna (2021) CS-DiD DRIPW estimator. Sample: manufacturing and market-services industries (EU15), staggered BIT entry 2004–2010. Standard errors via Mammen wild bootstrap (999 reps), resampled at the country–industry panel-unit level. Treatment timing is assigned at the country level, so identification rests on 15 countries; country-level clustering is reported as robustness in Appendix A.3. **VC TFP**: cumulative VC TFP growth (log-points). **Bus. intangibles**: log annual business intangible investment (flow). **Tech. intangibles**: log annual technology intangible investment (flow). $k = 0$ is the year the BIT enters into force; $k = -1$ is the baseline (omitted).

geneity, we partition the sample by the industry-level median of each industry's pre-treaty share of Chinese value added (TiVA share) and estimate the Callaway–Sant'Anna event study with doubly-robust inverse-propensity weighting (Callaway and Sant'Anna, 2021; Sant'Anna and Zhao, 2020) separately in the high- and low-exposure partitions.

Table 22 in Appendix A.3 reports the long-run (k=+7) ATTs and joint pre-trend χ^2 tests within each partition, together with a direct high-vs-low Wald difference test. Both partitions display positive long-run BIT effects for all four outcomes—VC TFP, KLEMS TFP, log total intangible stock, and log business intangible stock—and none of the high-vs-low differences is statistically distinguishable from zero at conventional levels. The largest partition differential is −6.3 log points for cumulative KLEMS TFP (p=0.29); the other three differentials are within ±3.5 log points of zero. We therefore do not find robust industry-level heterogeneity in the BIT effect along the China-exposure margin.

Pre-trend behaviour, however, is meaningfully heterogeneous, and it does not favour either partition uniformly. The high-exposure subsample rejects parallel trends for the business intangible stock ($p < 0.001$) and for cumulative KLEMS TFP ($p = 0.002$), where the low-exposure subsample does not. But the low-exposure subsample rejects parallel trends for cumulative VC TFP ($p = 0.011$)—our primary outcome—where the high-exposure subsample does not, and it is borderline for the total intangible stock ($p = 0.068$). The high-exposure pattern is consistent with a narrative in which industries already deeply integrated with Chinese supply chains were on a distinctive pre-BIT trajectory, so that part of the observed post-BIT movement reflects a pre-existing gap.

Together, the event-study and heterogeneity findings document a distinctive pattern: integration is associated with broad-based productivity gains that do not sort systematically along the China-exposure margin at conventional significance levels. The intangible investment response is present in both high- and low-exposure industries, though the high-exposure subsample carries a stronger pre-trend caveat that should temper the interpretation of its post-BIT ATT.

Elasticity estimates. The event study establishes that BITs are followed by an increase in VC TFP and in business-intangible investment. The long-run (t+7) ATT estimates already indicate a strong link, in that the two gains are of comparable magnitude. However, the BIT plainly operates through channels other than business-intangible investment—output, value added, and the organization of upstream production all respond—so the exclusion restriction that a causal reading would require does not hold. To assess how different forms of intangible investment relate to VC TFP growth more systematically, we turn to fixed-effects panel estimates. These estimates complement the event-study evidence. By conditioning on additional controls and exploiting within-chain variation over time, they provide a more systematic characterization of the intangible–TFP relationship and supply the key elasticity inputs for the structural model.

Table 4 reports the results. Column (1)—the level fixed-effects specification for the pre-2008 sample with the business–technology interaction—is the primary specification. The contemporaneous business-intangible coefficient is $\hat{\beta}_A = 0.356^{***}$ (SE = 0.100), while the lagged coefficient is -0.385^{***}, of approximately equal and opposite magnitude.[12] The interaction term $\hat{\beta}_{AZ} = 0.008$ is positive, indicating business–technology complementarity, though imprecisely estimated at conventional significance levels. The technology coefficient is small and insignificant ($\hat{\beta}_Z = -0.030$, SE = 0.062), confirming that the direct association between technology intangibles and VC TFP growth is weaker than the business-capability effect.

Column (2) shows the post-2008 results: $\hat{\beta}_A$ falls to 0.195^{*}, roughly half the pre-2008 figure—a regime shift that is central to the post-2008 asymmetry narrative. Column (3), a static system GMM, yields a point estimate close to the pre-2008 level-FE result ($\hat{\beta}_A = 0.380^{**}$, AR(2) $p = 0.206$); because its Hansen test is degenerate, we read it as a robustness check rather than as stronger evidence on causality. Column (4) uses conventional EU KLEMS TFP as the outcome rather than VC TFP; the business-intangible coefficient ($\hat{\beta}_A = 0.043$, SE = 0.525) is statistically indistinguishable from zero, confirming that the business-intangible relationship is substantially more clearly identified

[12]Consider $g_t = \beta_0 x_t + \beta_1 x_{t-1} + e_t$. If x rises permanently by Δx in period t, the impact on growth is $\beta_0 \Delta x$. From period $t+1$ onward the incremental growth contribution is approximately $(\beta_0 + \beta_1)\Delta x$; when $|\beta_0| \approx |\beta_1|$ the growth burst is temporary and the *level* shift is permanent—consistent with the model's comparative-statics interpretation.

using the value-chain productivity measure. Additional robustness specifications—no-cross and first-differenced variants—are in Appendix A.4 and yield $\hat{\beta}_A$ in the range 0.36–0.40 across all pre-2008 specifications. These elasticity estimates form the empirical inputs to the structural model calibration in Section 6.

6 A quantitative-theory framework

The empirical evidence documents that EU integration with China is associated with gains in value-chain productivity and business-intangible investment, that these gains are broad-based rather than concentrated in the most China-exposed industries, and that the business-intangible channel weakens sharply after 2008. These patterns are informative about European industries, but they leave two important questions open. First, the EU industry panel contains no direct information on Chinese productivity dynamics: we observe the European side of the integration episode but not the complementary Chinese response to the same shocks. Second, the reduced-form estimates are partial-equilibrium in nature—they quantify the average BIT effect for a European industry, but they do not capture the general-equilibrium feedbacks through trade volumes, relative prices, and scale reallocation that shape the aggregate outcome.

To address these limitations, we calibrate a stylized two-region general-equilibrium model of trade and intangibles. The model is deliberately parsimonious and is presented as a framework for organizing and extending the empirical findings, not as a complete structural account of the European productivity slowdown. The value-chain structure of the data is collapsed into Europe and a China/upstream counterpart so that the model can be used to interpret the empirical moments.

The model traces two distinct episodes through steady-state comparisons. A trade liberalization shock—representing the BIT-driven opening of EU–China goods trade—generates the initial integration gains: lower trade costs raise market access, reward investment in productive capabilities, and lift value-chain productivity on both sides. A second set of shocks to technology capacity and business capability then generates the divergent, asymmetric trajectories of the post-2008 period.

Table 4: Intangibles and TFP growth

	VC TFP			KLEMS TFP
	(1) Level FE w/cross 1995–2007	(2) Level FE w/cross 2008–2017	(3) GMM no cross 1995–2007	(4) Level FE w/cross 1995–2007
Bus. intangibles (β_A)	0.356***	0.195*	0.380**	0.043
	(0.100)	(0.109)	(0.162)	(0.525)
Bus. intangibles$_{t-1}$	−0.385***	−0.270**	−0.356**	−0.054
	(0.108)	(0.115)	(0.147)	(0.490)
Tech. intangibles (β_Z)	−0.030	0.066	0.177*	−0.423
	(0.062)	(0.092)	(0.094)	(0.447)
Tech. intangibles$_{t-1}$	0.013	−0.105	−0.179**	0.391
	(0.063)	(0.089)	(0.090)	(0.423)
Bus. × Tech. (β_{AZ})	0.008	−0.009		0.049
	(0.010)	(0.012)		(0.057)
Bus.$_{t-1}$ × Tech.$_{t-1}$	−0.004	0.015		−0.043
	(0.010)	(0.012)		(0.054)
Tangibles	−0.412***	−0.300***	−0.673***	−0.021
	(0.062)	(0.076)	(0.157)	(0.080)
Tangibles$_{t-1}$	0.391***	0.244***	0.675***	−0.029
	(0.067)	(0.080)	(0.159)	(0.089)
Share foreign mkt. serv.	−0.770***	−0.865**	1.185	−2.940***
	(0.226)	(0.369)	(1.205)	(0.378)
Share foreign mkt. serv.$_{t-1}$	1.044***	1.051**	−0.883	3.282***
	(0.236)	(0.414)	(1.152)	(0.516)
Observations	3173	3159	3173	2751
Number of groups	312	318	312	263
Within R^2	0.203	0.190		0.096
Instruments			31	
AR(2) (p)			0.206	
Hansen (p)			n.d.	

Notes: Columns (1), (2), and (4): fixed-effects estimates with year effects and robust standard errors clustered at the value-chain panel identifier; regressors are log intangible stocks and their lags. Column (3): two-step static system GMM (no lagged dependent variable); lag window [2,5], collapsed instrument matrix, 31 instruments. The dependent variable is annual VC TFP growth in columns (1)–(3); column (4) uses conventional EU KLEMS TFP growth to illustrate that the business-intangible relationship is substantially weaker with the conventional industry-level measure. AR(2): Arellano–Bond test for second-order serial correlation. Hansen test has negative degrees of freedom (degenerate); AR(2) is the primary diagnostic for column (3). Additional robustness specifications (no-cross and first-differenced variants) are in Appendix A.4. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Because China's TFP is substantially more elastic with respect to technology intangibles than Europe's, the same global technology developments benefit China disproportionately—producing the pattern of relative EU slowdown that the reduced-form evidence identifies.

The framework has two countries, Europe and China, each producing a differentiated manufactured good. Consumers in both countries combine domestic and imported goods through standard Armington-CES demand, so that lower trade costs increase market access and raise incentives to invest in productive capabilities.

Each country's manufacturing TFP depends on two *effective intangible bundles*. The business-intangible bundle (A_i) combines the country's own domestic organizational capability (a_i) with business-intangible knowledge that spills over from the other country. The technology-intangible bundle (Z_i) does the same for technology-related knowledge. Both bundles are attenuated by *diffusion barriers* (τ_A, τ_Z) that reduce the effective cross-border spillover. The country-specific parameters θ_i^A and θ_i^Z are the *elasticities of TFP with respect to each intangible bundle*: θ_i^A measures the percentage increase in country i's manufacturing productivity when its business-intangible bundle expands by one percent, and θ_i^Z does the same for the technology bundle. The calibration shows that these elasticities differ markedly between Europe and China—and the direction of this difference is the model's central result.

Domestic business capability a_i is produced from investment in business intangibles, and the payoff to that investment depends on the *delivered scale* of the region's tradable-goods production—the total quantity of manufactured output absorbed across domestic and export markets. The idea is that organizational capability is accumulated jointly with production experience and market presence (Prescott and Visscher, 1980; Atkeson and Kehoe, 2005). The coefficient μ_A governs the interaction between the intangible stock and scale. It is a complementarity, not a scale elasticity, and its content has to be read off the calibrated equilibrium: at Step 1 the marginal elasticity of capability with respect to the stock, $\eta_A + \mu_A \ln S_i$, is 1.10 in Europe and 0.56 in China, so the interaction opens a wedge in the capability payoff between a large region and a small one. It is not a scale economy in the usual sense. Because (11) is written in the product of two logarithms, the derivative with respect to scale, $\mu_A \ln K_i^A$, is -0.32 in Europe and -1.01 in China at the calibrated stocks. This is why Chinese capability falls over the liberalization step even as Chinese intangible investment rises sharply, and it is why we treat μ_A as a fitted feature of the capability block rather than as a structural parameter.

Relation to previous models. Each block is stylized. The trade block is a standard CES structure with an iceberg friction, adopted for tractability: there is no free entry, firm heterogeneity or selection margin, so the model cannot speak to reallocation across firms. Lee and Yi (2018) is the closest global value-chain benchmark, in which lower trade costs have magnified aggregate effects because production crosses several linked stages; we abstract from their stage-by-stage specialization and worker sorting so as to concentrate on capability accumulation and cross-border diffusion. The TFP and capability blocks draw on the intangible-capital literature (Corrado et al., 2017; Brynjolfsson et al., 2021; van Ark et al., 2024), and the capability block is closer to the organization-capital mechanisms of Prescott and Visscher (1980) and Atkeson and Kehoe (2005) than to scale economies in the trade sense. The cross-border bundles are a reduced form for international knowledge diffusion, which Eaton and Kortum (1999), Keller (2004) and Buera and Oberfield (2020) show is neither automatic nor symmetric across countries; the region-specific θ_i^Z summarize that heterogeneity in absorptive capacity.

Production. Each country $i \in \{E, C\}$ has a Cobb-Douglas manufacturing production function,

$$q_i = T_i K_i^{\alpha} L_i^{1-\alpha}, \tag{7}$$

where q_i is manufacturing output, T_i is TFP, K_i is physical capital, L_i is a fixed labour endowment, and α is the capital share.

TFP. Manufacturing TFP depends on a size shifter $\bar{T}_i$ and on effective bundles of business intangibles A_i and technology intangibles Z_i:

$$T_i = \bar{T}_i \, A_i^{\theta_{A,i}} \, Z_i^{\theta_{Z,i}}, \tag{8}$$

with region-specific elasticities $\theta_{A,i}$ and $\theta_{Z,i}$. These elasticities are the model's key new element relative to Lee and Yi (2018): they are what allow the two regions to respond asymmetrically to the same technology improvement.

Intangible bundles. Each country's TFP draws on both its own domestic capability and knowledge that spills over from the other country, attenuated by diffusion barriers τ_A and τ_Z. The effective business-intangible bundle for Europe is

$$A_E = \left[a_E^{\frac{\sigma_A - 1}{\sigma_A}} + \left(\frac{a_C}{\tau_A} \right)^{\frac{\sigma_A - 1}{\sigma_A}} \right]^{\frac{\sigma_A}{\sigma_A - 1}}, \tag{9}$$

where a_i is domestic business capability and σ_A is the substitution elasticity between domestic capability and cross-border spillovers. The bundles use equal unit weights (symmetric CES), so diffusion barriers τ_A and τ_Z are the sole source of home bias. The technology bundle is defined analogously,

$$Z_E = \left[z_E^{\frac{\sigma_Z - 1}{\sigma_Z}} + \left(\frac{z_C}{\tau_Z} \right)^{\frac{\sigma_Z - 1}{\sigma_Z}} \right]^{\frac{\sigma_Z}{\sigma_Z - 1}}, \tag{10}$$

where z_i is country i's domestic technology-service capacity, and the corresponding expression for China is obtained by swapping E and C in (10). In each steady state, z_i is treated as an exogenous scalar: it sets the baseline level of the country's technology endowment for that step of the model, independently of current investment flows. Europe's technology capacity is normalized to $z_E = 1$ throughout. China's technology capacity z_C takes the same value in Steps 0 and 1; Step-2 comparative statics perturb it around the Step-1 equilibrium. Its baseline value is reported in Table 5 and appears in the shock tables of Table 8 and Appendix A.6.

Capability formation and scale. Domestic business capability a_i is produced from the *stock* of business intangibles K_i^A and delivered production scale S_i via a log-linear relationship:

$$\ln a_i = \ln \bar{a}_i + \eta_A \ln K_i^A + \mu_A (\ln K_i^A)(\ln S_i), \tag{11}$$

where $\bar{a}_i$ is an exogenous capability shifter, η_A is the direct investment elasticity (normalized to 1), $\mu_A \geq 0$ is the stock–scale interaction coefficient[13], while gross investment is a fixed share of final output and the stock follows from the steady-state accumulation identity $I_i^A = \delta_A K_i^A$:

$$I_i^A = s_i^A Y_i, \qquad K_i^A = I_i^A / \delta_A, \tag{12}$$

where δ_A is the depreciation rate of business intangibles. Building capability on the stock rather than on the flow matters for the mapping to the data: the panel elasticities in Section 5 are estimated on log intangible *stocks*, so the regressor in the model's moment counterpart is now the same object as the regressor in the regression. Delivered scale is the total quantity of domestic manufactured goods absorbed across both markets:

$$S_i \equiv y_{ii} + y_{ij}, \tag{13}$$

where y_{ii} is domestic absorption and y_{ij} is the quantity delivered to the other country.

Trade and market clearing. Consumers in each country have Armington-CES preferences over domestic and imported goods with elasticity σ_M. The price index dual to the CES aggregator is

$$P_i = \left[p_i^{1-\sigma_M} + \left(\tau_M p_j\right)^{1-\sigma_M} \right]^{\frac{1}{1-\sigma_M}}, \tag{14}$$

[13]Because μ_A multiplies the product of two logarithms, it is a complementarity coefficient rather than an elasticity: the marginal elasticity of capability with respect to the stock is $\eta_A + \mu_A \ln S_i$ and that with respect to scale is $\mu_A \ln K_i^A$. Neither is invariant to the units in which K_i^A and S_i are measured. Expressing both arguments relative to explicit reference levels would make μ_A normalization-invariant; we have not done so, and Appendix A.7 reports what the interaction does at the calibrated values.

where p_i is the factory-gate price of country i's good and τ_M is the iceberg goods trade cost. Market clearing for manufacturing output requires

$$q_i = y_{ii} + \tau_M \, y_{ji}, \tag{15}$$

where y_{ji} is the quantity of country i's good absorbed in country j, and the factor τ_M accounts for the iceberg units melted in transit. This two-region Armington-CES specification with an iceberg friction is a standard device in quantitative trade work, adopted here for tractability. There is no free entry, no firm-level heterogeneity, and no selection margin, so the model cannot speak to reallocation across firms.

Capital market equilibrium. In steady state the return to capital satisfies

$$\frac{r_i}{P_i} = \bar{R} \equiv \frac{1}{\beta} - (1 - \delta_K), \tag{16}$$

where β is the discount factor and δ_K is the depreciation rate. The equilibrium rental rate equals the marginal product of capital,

$$r_i = \alpha \, p_i \, T_i \, K_i^{\alpha-1} \, L_i^{1-\alpha}. \tag{17}$$

A steady-state equilibrium is a vector of prices and quantities such that firms maximize profits, households maximize utility, and all goods and factor markets clear at the rate $\bar{R}$.

6.1 Calibration

The model is parameterized through steady-state comparisons across three steps. **Step 0** is the high-trade-cost benchmark prior to liberalization. **Step 1** reduces goods trade costs to represent the BIT-driven liberalization phase; this is the primary shock we use to generate the pre-2008 integration dynamics. **Step 2** perturbs the productivity shifters and diffusion barriers around the Step 1 equilibrium, to characterize the disturbances that could account for the weaker European trajectory after 2008. The shifters $\bar{T}_i$, $\bar{a}_i$ and z_i set each region's baseline productivity, organizational capacity and technology capacity independently of current investment flows, so shocks to them represent changes in structural advantage rather than in measured investment.

Table 5 reports the full parameterization. Standard macro and trade parameters are set to conventional values. The substitution elasticities $\sigma_A = \sigma_Z = 6$ govern the degree to which domestic capability and cross-border spillovers substitute for each other in the intangible bundles. The investment elasticity η_A is normalized to one, so that a one-percent increase in intangible investment raises domestic capability by one percent before the scale interaction is applied.

We discipline the model with moments drawn from the pre-2008 empirical analysis, the period for which identification is cleanest. The moments connect the two empirical exercises directly.

The first four moments capture the BIT liberalization episode and are anchored to the event-study ATT estimates at the seven-year horizon. We use the t+7 ATT as a *long-horizon calibration anchor* for the steady-state change implied by the BIT shock. This is the longest horizon at which all treated countries still contribute to the estimate, but it should not be read as evidence that a new steady state has been reached: the post-BIT trajectories are still rising at $k = +7$, which means the t+7 ATT is if anything a lower bound on the eventual response. Specifically, the European VC TFP response in Step 0→1 is set to $\hat{\tau}_{+7}^{TFP} = 0.102$ (from Table 3), the European business-capability response is set to $\hat{\tau}_{+7}^{bus} = 0.105$ (log business intangible investment at $k = +7$, converted to a stock change), the upstream/China TFP response is set to 0.030, and the pre-liberalization EU/China nominal GDP ratio is set to 3.0. Our data contain no Chinese productivity panel, and 0.030 is a round value chosen to represent a positive but substantially smaller upstream response than Europe's, consistent with the upstream component of the VC TFP decomposition in Section 4.1.

The remaining three moments come from column (1) of Table 4 (level FE pre-2008 with cross-elasticity): the contemporaneous elasticities $\hat{\beta}_A = 0.356$ (TFP response to business intangibles), $\hat{\beta}_Z = -0.030$ (TFP response to technology intangibles), and $\hat{\beta}_{AZ} = 0.008$ (the interaction between business and technology intangible investment). We use the contemporaneous ($t = 0$) regression coefficient as the calibration target because the model has TFP *growth* as the dependent variable.

The t-period estimate gives the first-period impact of a permanent increase in intangible investment. Because the lagged coefficient is of approximately equal and opposite magnitude, the level of TFP is permanently shifted while the growth rate returns to steady state—consistent with the model's comparative-statics interpretation. The estimates are consistently similar across specifications: the static GMM (column (3) of Table 4) and the additional no-cross and first-differenced variants in Appendix A.4 all yield $\hat{\beta}_A$ in the range 0.36–0.40. This range, rather than the ratio of the two event-study responses, is what we target.

Estimation by minimum distance. Parameters are estimated by minimizing a weighted quadratic loss in relative deviations between model moments and empirical targets:

$$\min_{\Theta} \sum_{k} \omega_k \left(\frac{m_k(\Theta) - \widehat{m}_k}{0.5|m_k(\Theta)| + 0.5|\widehat{m}_k| + \epsilon} \right)^2, \tag{18}$$

where ω_k are moment weights and ϵ prevents division by zero when a target is close to zero; we use equal weights and $\epsilon = 10^{-8}$. We minimize by Nelder–Mead from a spread of starting values and report the best, at which the objective is 0.386.

Because β_A is the one moment the model cannot match, its relative deviation accounts for most of the objective at the optimum and the remaining six are matched closely along a ridge of parameter combinations. The objective is correspondingly flat in the technology elasticities: solutions with θ_E^Z between 0.037 and 0.086 lie within 0.5 % of one another. We report the best-fitting point and read the *difference* $\theta_C^Z - \theta_E^Z \approx 0.50$, rather than the ratio, as the quantity the moments pin down. One parameter is genuinely not identified by curvature: z_C reaches whatever upper bound is imposed on it, so we predetermine it rather than treat it as a calibrated technology level, and report sensitivity across its range (Appendix A.7).

Appendix A.7 reports two further calibrations: Spec B, which excludes β_{AZ} from the objective and targets a technology coefficient of the opposite sign, and Spec C, which fixes $\mu_A = 0$. Spec B still returns $\theta_C^Z \gg \theta_E^Z$, so the sign of the asymmetry is not an artifact of targeting the cross-elasticity; Spec C does not converge to an acceptable fit at all, which tells us that the interaction term is needed for the model to reproduce the cross-region size relationship but says nothing directly about the asymmetry. Appendix A.5 sets out which moment moves which parameter and discusses why z_C is treated as a predetermined parameter. Appendix A.7 shows that constraining θ_E^A at lower values (0.40–0.80) actually increases the θ_C^Z/θ_E^Z ratio, so the asymmetry is not an artifact of the particular θ_E^A chosen by the optimizer. Appendix A.7 reports re-calibrations across a range of values for the predetermined technology capacity z_C and for the diffusion barriers.

Table 5: Macroeconomic model calibration: parameters

	Europe	*China*
Panel A. Estimated parameters		
Elasticity in business intangibles (θ_i^A)	0.964	0.351
Elasticity in tech intangibles (θ_i^Z)	0.037	0.540
Stock–scale interaction (μ_A)	0.384	
Size shifter ($\bar{T}_i$)	1.228	0.093
Panel B. Predetermined parameters		
China technology level (z_C)	1.000	
China business shifter ($\bar{a}^C$)	0.650	
Discount factor (β)	0.96	
Depreciation rate (δ_K)	0.06	
Armington elasticity (σ_M)	4.0	
Substitution elasticities ((σ_A, σ_Z))	6.0	6.0
Capital share (α)	0.35	
Labor endowments (L_E, L_C)	1.0	4.0
Business investment shares ((s_E^A, s_C^A))	0.06	0.03
EU business shifter ($\bar{a}_E$)	1.0	
EU technology shifter (z_E)	1.0	
Stock elasticity (η_A)	1.0	
Depreciation, business intangibles (δ_A)	0.20	
Trade costs, Step 0 (τ_M^H)	1.60	
Trade costs, Step 1 (τ_M^L)	1.10	
Diffusion barriers (τ_A, τ_Z)	1.20, 1.10	

Notes: Panel A lists the seven estimated parameters; Panel B lists the predetermined values, whose basis is discussed in Appendix A.5. Targets: $\hat{\beta}_A = 0.356$, $\hat{\beta}_Z = -0.030$, $\hat{\beta}_{AZ} = +0.008$ from column (1) of Table 4 (level FE pre-2008); event-study anchors $\hat{\tau}_{+7}^{TFP} = 0.102$ and $\hat{\tau}_{+7}^{bus} = 0.105$ from Table 3; parameter bounds $\theta_{\max} = 2.00$, $\bar{T}_{\max} = 10.00$. All four estimated elasticities are interior solutions; the technology asymmetry $\theta_C^Z \gg \theta_E^Z$ emerges from matching the empirical moments. z_C is predetermined at $1.000 = z_E$, so baseline technology capacity is symmetric across regions and every cross-region asymmetry in the calibration is carried by the estimated conversion elasticities. When z_C is instead included as an estimated parameter it runs to whatever upper bound is imposed, so it is not identified by the moments; sensitivity across $z_C \in [0.5, 5]$ is reported in Appendix A.7. Robustness calibrations Spec B and Spec C are reported in Appendix A.7.

Table 6: Matched empirical moments and model counterparts

Moment	Primary calibration Target	Model	Gap
Pre-lib. EU/CN GDP ratio	3.000	3.055	+0.055
EU VC TFP response ($\hat{\tau}^{TFP}_{+7}$)	0.102	0.101	–0.001
Upstream/China TFP response	0.030	0.030	0.000
EU business-capability response	0.105	0.105	0.000
Business elasticity (β_A)	0.356	0.676	+0.320
Technology elasticity (β_Z)	–0.030	–0.030	0.000
Interaction elasticity (β_{AZ})	0.008	0.008	0.000

Notes: Targets from column (1) of Table 4 (level FE pre-2008): $\hat{\beta}_A = 0.356$, $\hat{\beta}_Z = -0.030$, $\hat{\beta}_{AZ} = +0.008$; event-study anchors from Table 3. The β_A gap is a target miss, not a matched moment; see Section 6.2. The other six moments are matched closely. Robustness calibrations are in Appendix A.7.

6.2 Results and interpretation

Pre-2008 liberalization gains

We begin by studying the global value chain liberalization shock. In our model, this episode boils down to one variable: the goods trade cost τ_M (Table 5, Panel B). A reduction in τ_M from $\tau_M^H = 1.60$ to $\tau_M^L = 1.10$ represents the BIT-driven liberalization phase, expanding market access for both regions and raising output, the delivered scale of production, and—through (11)—business-intangible investment and capability in Europe.

Table 7 reports the model steady states before and after the decline in goods trade costs. The only structural change between Step 0 and Step 1 is the reduction in τ_M. Despite its simplicity, the model generates a joint increase in productivity and output in both regions. The direct match to the empirical liberalization episode is visible in the targeted objects: European manufacturing TFP rises by 10.1%, upstream/Chinese manufacturing TFP rises by 3.0%, and Europe's effective business-capability bundle A_E rises by 10.5%.

The same liberalization step also generates informative non-targeted moments. Europe's domestic business-capability component a_E rises by 18.0%, European final output Y_E by 25.0%, and European manufacturing output q_E by 19.1%, while the effective technology bundle Z_E is unchanged by construction. China's domestic capability component a_C falls (−5.5%) even though the Chinese business-intangible stock rises by 47%; this is a consequence of the interaction term in (11), discussed below. These patterns line up well with the reduced-form evidence in Section 5. The BIT event-study results show a clear increase in VC TFP, an increase in output and value added, and a positive response in business-related intangibles, whereas technology-related intangibles display much weaker movement (Figures 1, 3, and 4). In the model, A_E is the closest analogue to the capability bundle that enters productivity directly, while a_E is the closer analogue to a purely domestic business-capability component. The fact that both increase, and that a_E rises somewhat more strongly than A_E, helps bracket the empirical business-intangible response rather than tie it to a single model object.[14]

Table 7: Steady states around the integration shock (Step 0 vs. Step 1)

	Step 0	Step 1	% change (0→1)
Goods trade cost (τ_M)	1.60	1.10	–
EU manufacturing TFP (T_E)	0.684	0.753	10.11
China manufacturing TFP (T_C)	0.112	0.115	3.04
EU effective business bundle (A_E)	0.529	0.584	10.51
EU domestic business capability (a_E)	0.336	0.397	17.95
China domestic business capability (a_C)	0.158	0.149	−5.48
EU effective tech bundle (Z_E)	2.193	2.193	0.00
EU final output (Y_E)	1.149	1.436	24.98
China final output (Y_C)	0.327	0.482	47.20
EU manufacturing output (q_E)	1.106	1.317	19.05
China manufacturing output (q_C)	0.287	0.338	17.96
EU business-intangible stock (K_E^A)	0.345	0.431	24.98
EU delivered scale (S_E)	1.061	1.289	21.48
China delivered scale (S_C)	0.251	0.318	26.60

Notes: Step 0 is the pre-liberalization steady state with high goods trade costs and Step 1 is the post-liberalization steady state after the fall in goods trade costs. Values are in model units and levels are rounded for display; the percentage changes in the final column are computed from the unrounded model output and therefore need not reproduce exactly from the displayed levels. Delivered scale S_i is the total quantity of region i's manufactured goods absorbed across both markets (domestic and export destination). A_E is the effective business-intangible bundle combining domestic capability a_E and cross-border business services; a_E is the purely domestic component. The percentage changes are the main objects used in the comparison to the BIT evidence.

[14] The event-study evidence identifies the direction and broad magnitude of the liberalization response more cleanly than the exact mapping from empirical business-intangible measures to a single model object. For that reason, we use A_E as the calibration target because it enters the productivity block directly, while also reporting a_E as a non-targeted comparison.

Quality of fit and interpretation. Table 6 shows that the calibration matches six of the seven targeted moments closely—the two event-study anchors, the pre-liberalization GDP ratio, the Chinese TFP response, and both the technology elasticity β_Z and the interaction elasticity β_{AZ}—but that it does not reproduce the panel estimate of the business-intangible elasticity. The model implies $\hat{\beta}_A \approx 0.68$ against a target of 0.356, a gap of +0.319. The model cannot deliver β_A below about 0.68: general-equilibrium feedback through output, investment and delivered scale ties TFP to the business bundle more tightly than the reduced-form panel relationship does. We read this as a failure of the capability block (11), not as evidence against the panel estimate, and the asymmetry results below turn on the directions of comparative statics rather than on this level.

The main result of the calibration is the technology asymmetry. Europe's technology-to-TFP elasticity ($\theta_E^Z = 0.037$) is substantially smaller than China's ($\theta_C^Z = 0.540$). The asymmetry is not imposed: it is what the model requires in order to match the empirical moments, and in particular the negative β_Z. The mechanism is that when $\theta_C^Z \gg \theta_E^Z$, a rise in European technology intangibles raises China's TFP disproportionately through the cross-border bundle, which depresses Europe's relative position and delivers a negative contemporaneous coefficient on European technology intangibles.

One qualification belongs with this. Since there is no Chinese productivity panel in our analysis, θ_C^Z is inferred from European moments through the model's general-equilibrium structure rather than measured. Appendix A.7 reports the robustness analysis: the gap survives dropping β_{AZ} from the objective and re-calibrating across ranges of z_C, the diffusion barriers and θ_E^A.

Europe has the stronger business-capability elasticity ($\theta_E^A = 0.964$ versus $\theta_C^A = 0.351$), which is consistent with organizational capital being Europe's comparative advantage in the pre-2008 period—subject to the caveat that β_A is the one moment the calibration does not match.

The stock–scale interaction plays a supporting role, by allowing the capability–stock relationship to differ between Europe and China. Consistent with this, Appendix A.7 shows that fixing $\mu_A = 0$ (Spec C) makes the calibration fail on the pre-liberalization EU/China GDP ratio. The primary driver of the asymmetry—the gap in technology elasticities $\theta_C^Z \gg \theta_E^Z$—is nevertheless robust to this restriction.

External comparison. Two calibrated objects can be placed loosely against the literature. The integration friction τ_M summarizes legal uncertainty, investment barriers, contract enforcement and goods-market access jointly rather than a measured trade cost; its fall from 1.60 to 1.10 is a 31 % reduction in the delivered-price markup, of the same order as the decline in bilateral trade costs that Novy (2013) infers for the United States between 1970 and 2000. The *direction* of the technology-elasticity gap is what the catch-up literature predicts: Griffith et al. (2004) find that industries further from the frontier draw more productivity growth from R&D, and Hu and Jefferson (2004) find substantial returns to R&D in Chinese industry. We do not compare magnitudes, since θ_i^Z is the exponent on a CES bundle in model units rather than an R&D elasticity.

Interpreting the post-2008 divergence

Having established that the model replicates the pre-2008 liberalization dynamics, we now use it to characterize what types of structural shocks can account for the post-2008 divergence between Europe and China. The exercise asks: given the calibrated parameters, which fundamental shifters generate outcomes consistent with the observed productivity slowdown in Europe and the continued gains in China?

The model's most useful output here is a statement about the *direction* of shocks consistent with the post-2008 pattern. It is an interpretation rather than a test: we do not ask the model to reproduce the measured post-2008 attenuation of $\hat{\beta}_A$, and doing so would be a natural extension. Table 8 reports this by computing how key endogenous variables respond to four 1 % increases in the fundamental productivity shifters around the calibrated Step 1 steady state. Each experiment raises one exogenous shifter by 1 % while holding all others fixed at their Step 1 values, so the resulting changes in TFP, output, and capability isolate the direct and general-equilibrium effects of that single source of improvement. The four experiments cover Europe's business-capability shifter ($\bar{a}_{E,1}$, the

Table 8: Asymmetric output and TFP responses to 1% productivity shocks (around Step 1)

	$\bar{a}_{E,1}$+1%	$z_{E,1}$+1%	$\bar{a}_{C,1}$+1%	$z_{C,1}$+1%
	EU bus. cap.	EU technology	China bus. cap.	China technology
EU TFP (T_E)	2.911	−0.044	1.005	−0.060
China TFP (T_C)	0.944	0.232	0.358	0.248
EU output (Y_E)	4.289	−0.037	1.478	−0.057
China output (Y_C)	2.090	0.268	0.760	0.283
EU business capability (A_E)	3.022	−0.066	1.043	−0.080
China business capability (A_C)	2.711	−0.076	1.023	−0.090

Notes: Entries are percentage changes relative to the Step 1 steady state after a 1% increase in the indicated shifter. Values from the primary calibration ($\mu_A = 0.384$, $\theta^Z_E = 0.037$, $\theta^Z_C = 0.540$). Note that both technology experiments *lower* EU TFP slightly: with θ^Z_E small, Europe's direct gain from a larger technology bundle is outweighed by the general-equilibrium loss of market share to China. Additional perturbation results for trade frictions, diffusion barriers, and elasticity changes are reported in Appendix A.6.

exogenous baseline level of Europe's organizational and business-service production), Europe's technology shifter ($z_{E,1}$), China's business-capability shifter ($\bar{a}_{C,2}$), and China's technology shifter ($z_{C,2}$).

Three patterns in Table 8 characterize the core asymmetry.

Business capability is the primary EU lever. A 1 % increase in Europe's business-capability shifter raises EU TFP by 2.91 % and EU output by 4.29 %. The multiplier of roughly four is the model's most powerful internal mechanism: higher capability raises TFP, which raises output and delivered scale, which raises intangible investment and capability again. It is why the business channel dominates the technology channel in every experiment, and it is the model's counterpart of the scale-and- diffusion argument in Draghi (2024). Spillovers run across the value chain: China's TFP and output also rise (0.94 % and 2.09 % respectively), consistent with integrated production. A 1 % shock to China's business capability similarly benefits Europe (EU TFP +1.01 %, EU output +1.48 %). Business-capability shocks, whether in Europe or China, expand both regions' production, but the domestic effect is substantially larger than the spillover.

Technology shocks are asymmetric. A 1 % increase in Europe's technology shifter raises China's TFP by 0.23 % and China's output by 0.27 %, while EU TFP *falls* by 0.04 %. In the calibrated model, that is, a European technology improvement raises Chinese productivity and lowers European productivity.

The arithmetic is worth setting out, because the result rests on two offsetting terms of similar size. Baseline technology capacity is symmetric ($z_C = z_E = 1$), so the European improvement raises the two effective bundles by almost the same amount: Z_E by 0.52 % and Z_C by 0.48 %. Europe's direct gain is therefore $\theta_E^Z \times 0.52 \approx 0.02$ percentage points. The realized change is -0.04, so the general-equilibrium term—the loss of market share as China becomes more competitive—is about -0.06 percentage points. Europe's net loss is small in absolute terms, so the claim we emphasize is the *contrast* between the two regions—China gains 0.23 where Europe gains nothing—rather than the precise magnitude for Europe. Appendix A.7 shows the pattern holds across the range of z_C and of the diffusion barriers that we examine.

A 1 % increase in China's own technology shifter works the same way: China's TFP rises by 0.25 % and EU TFP falls by 0.06 %.

The mechanism is the gap in technology elasticities. The calibration assigns a substantially higher elasticity of manufacturing TFP with respect to the effective technology-intangible bundle in China than in Europe: $\theta_C^Z = 0.540$ versus $\theta_E^Z = 0.037$, a difference of about 0.50. A one-percent increase in the same effective technology bundle therefore produces a larger direct TFP response in China than in Europe. General equilibrium then adds a second, separate step: the resulting shift in output changes delivered market scale and, through the capability block (11), feeds back into business capability. That feedback works *against* Europe rather than for it: in both technology experiments in Table 8 the business-capability bundles fall in both regions (-0.066 and -0.076 for an EU technology shock, -0.080 and -0.090 for a Chinese one), because technology-driven reallocation shifts delivered scale toward China and the capability block responds to scale. The elasticity gap creates the asymmetry in TFP, and the capability channel deepens rather than offsets it.

A reading of the post-2008 divergence. The asymmetry in Table 8 translates directly into a narrative for the post-2008 period. If, after 2008, technology diffusion runs disproportionately toward China—through global sourcing networks, knowledge spillovers, and China's rapid accumulation of technology-related intangibles—then the model predicts that China converts these gains into higher upstream productivity and a larger production scale (a 1 % Chinese technology shock raises Chinese TFP by 0.25 % while lowering EU TFP by 0.06 %), while Europe's value-chain productivity falls relative to China and Europe's upstream contribution remains comparatively subdued. This is consistent with the empirical finding that business intangibles are robustly associated with pre-2008 VC TFP growth but much less so after 2008, and that VC TFP weakens both in final industries and upstream. The robustness of this asymmetry prediction—across alternative calibrations and parameter restrictions—is verified in Appendix A.7.

7 Concluding remarks

European productivity growth has slowed markedly since the Global Financial Crisis, despite rapid advances in digital technologies and substantial investment in intangible assets. This paper revisits the slowdown through the lens of global value chains by constructing a panel of value-chain TFP (VC TFP) for EU15 industries over 1995–2017. The value-chain perspective matters because productivity and input growth increasingly reflect inter-industry linkages and cross-border production networks rather than isolated sectoral performance.

Our first contribution is measurement and decomposition. We aggregate primary inputs across the full chain to obtain VC TFP indexed by the final producer, and we decompose VC productivity into components originating in final industries and in upstream segments. The post-2008 slowdown is broad-based in this decomposition: VC TFP weakens both in final producer industries and in the upstream part of the chain. This shifts attention from a purely "within-industry" productivity story toward mechanisms that operate through value-chain organization, intermediate input use, and the productivity of upstream suppliers.

Our core empirical contribution is an account of how integration with China, mediated through intangible investment, shaped VC productivity before and after 2008. The event-study design exploits the staggered timing of bilateral investment treaties with China as a quasi-experiment. BITs are followed by cumulative VC TFP gains of 10.2% and business-related intangible investment gains of 10.5% at the seven-year horizon. These average effects are broad-based across the high- and low-China-exposure partitions: both groups show positive long-run responses, and their differences are not statistically distinguishable from zero.

Different panel estimation models provide complementary evidence on the underlying elasticities. Business-related intangible investment is strongly and positively associated with VC TFP growth prior to 2008 ($\hat{\beta}_A = 0.356^{***}$ in the primary specification), whereas this relationship falls by roughly half after 2008. Technology-related intangibles have a weaker direct effect in baseline specifications. Together, the event-study and panel evidence point to business-oriented intangible capabilities—organizational capital, management practices, branding, and related assets—as an important channel through which integration and scale translate into value-chain productivity improvements, and to a post-2008 regime in which these channels deliver smaller measured gains.

To organize and interpret these empirical patterns, we calibrate a stylized two-region model of trade and intangibles in which each region's manufacturing TFP depends on effective bundles of business and technology intangibles with region-specific elasticities. The contribution is to show that calibrating to the empirical moments requires substantially different technology-absorption elasticities for Europe and China ($\theta^Z_E = 0.037$ versus $\theta^Z_C = 0.540$), and that this elasticity gap predicts asymmetric effects of technology shocks: China converts technology improvements into proportionately larger TFP and output gains than Europe. The stock–scale interaction in the capability block contributes to the model's quantitative coherence, but the asymmetry is driven by the elasticity gap rather than by scale.

Taken together, the results suggest a two-stage interpretation of Europe's productivity experience. Before 2008, deeper integration was accompanied by higher value-chain productivity and greater investment in business intangibles, consistent with European firms using accumulated organizational capabilities to exploit larger and more international production networks. The post-2008 slowdown is broader: productivity weakens both in final industries and upstream, and the empirical association between business intangibles and VC TFP becomes substantially smaller. The quantitative framework shows how these facts can coexist if technology-related knowledge is converted into productivity more effectively outside Europe.

The model therefore shifts the question from whether integration generates gains to whether European firms possess the technological absorption capabilities needed to sustain them. A natural avenue for future work is to analyze richer sources of heterogeneity, such as firm-level organization and multinational network structure, and to quantify how changes in trade, technology, and geopolitically motivated frictions alter the productivity returns to intangible investment.

References

Acemoglu, D. and Tahbaz-Salehi, A. (2020). Firms, failures, and fluctuations: The macroeconomics of supply chain disruptions. *NBER Working Paper*, 27565.

Atkeson, A. and Kehoe, P. J. (2005). Modeling and measuring organization capital. *Journal of Political Economy*, 113(5):1026–1053.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–2168.

Autor, D. H., Dorn, D., and Hanson, G. H. (2016). The china shock: Learning from labor-market adjustment to large changes in trade. *Annual Review of Economics*, 8(1):205–240.

Baldwin, R. and Robert-Nicoud, F. (2014). Trade-in-goods and trade-in-tasks: An integrating framework. *Journal of International Economics*, 92(1):51–62.

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203.

Becker, S. O., Ekholm, K., and Muendler, M.-A. (2013). Offshoring and the onshore composition of tasks and skills. *Journal of International Economics*, 90(1):91–106.

Bloom, N., Sadun, R., and Van Reenen, J. (2012). Americans do IT better: US multinationals and the productivity miracle. *American Economic Review*, 102(1):167–201.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tohoku earthquake. *The Review of Economics and Statistics*, 101(1):60–75.

Bonadio, B., Huo, Z., Levchenko, A., and Pandalai-Nayar, N. (2021). Global supply chains in the pandemic. *Journal of International Economics*, 133:103534.

Bontadini, F., Corrado, C., Haskel, J., Iommi, M., and Jona-Lasinio, C. (2023). EUKLEMS & INTANProd: Industry productivity accounts with intangibles. Sources of growth and productivity trends: Methods and main measurement challenges. Technical report, EUKLEMS & INTANProd (Deliverable D2.3.1, Procurement procedure ECFIN/2020/O.P./0001).

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, 91(6):3253–3285.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372. Earlier version circulated as NBER Working Paper 25148 (2018, revised 2020).

Buera, F. J. and Oberfield, E. (2020). The global diffusion of ideas. *Econometrica*, 88(1):83–114.

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

Carvalho, V. M., Nirei, M., Saito, Y., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: evidence from the great east japan earthquake. *Quarterly Journal of Economics*, 136(2):1255–1321.

Chor, D., Manova, K., and Yu, Z. (2021). Growing like china: Firm performance and global production line position. *Journal of International Economics*, 130:103445.

Constantinescu, C., Mattoo, A., and Ruta, M. (2019). Does vertical specialisation increase productivity? *The World Economy*, 42(8):2385–2402.

Corrado, C., Haskel, J., and Jona-Lasinio, C. (2017). Knowledge spillovers, ict and productivity growth. *Oxford Bulletin of Economics and Statistics*, 79(4):592–618.

Corrado, C., Haskel, J., Jona-Lasinio, C., and Iommi, M. (2022). Intangible capital and modern economies. *Journal of Economic Perspectives*, 36(3):3–28.

Cuñat, A. and Zymek, R. (2025). Unbalanced trade 2.0. IMF Working Paper WP/25/191, International Monetary Fund.

Deng, M. and Liu, C. (2024). Sovereign risk and intangible investment. *Journal of International Economics*, 152:104009.

Draghi, M. (2024). The future of european competitiveness: In-depth analysis and recommendations. Technical report, European Commission. Report requested by the President of the European Commission.

Eaton, J. and Kortum, S. (1999). International technology diffusion: Theory and measurement. *International Economic Review*, 40(3):537–570.

Elliott, M., Golub, B., and Leduc, M. (2022). Supply network formation and fragility. *American Economic Review*, 112(8):2701–2747.

Fan, H., Li, Y. A., and Yeaple, S. R. (2018). On the relationship between quality and productivity: Evidence from china's accession to the wto. *Journal of International Economics*, 110:28–49.

Goldin, I., Koutroumpis, P., Lafond, F., and Winkler, J. (2024). Why is productivity slowing down? *Journal of Economic Literature*, 62(1):196–268.

Griffith, R., Redding, S., and Van Reenen, J. (2004). Mapping the two faces of R&D: Productivity growth in a panel of OECD industries. *Review of Economics and Statistics*, 86(4):883–895.

Guzman, J., Murray, F., Stern, S., and Williams, H. L. (2024). Accelerating innovation ecosystems: The promise and challenges of regional innovation engines. In *Entrepreneurship and Innovation Policy and the Economy*, volume 3, pages 9–75. University of Chicago Press. Earlier version circulated as NBER Working Paper 31541 (2023).

Houseman, S. N. and Mandel, M., editors (2015). *Measuring Globalization: Better Trade Statistics for Better Policy*. Books from Upjohn Press. W.E. Upjohn Institute, Kalamazoo, MI.

Hu, A. G. and Jefferson, G. H. (2004). Returns to research and development in Chinese industry: Evidence from state-owned enterprises in Beijing. *China Economic Review*, 15(1):86–107.

Jona-Lasinio, C. and Meliciani, V. (2019). Global value chains and productivity growth in advanced economies: Does intangible capital matter? *International Productivity Monitor*, 36:53–78.

Juhász, R., Lane, N., and Rodrik, D. (2024). The new economics of industrial policy. *Annual Review of Economics*, 16:213–242. Earlier version circulated as NBER Working Paper 31538 (2023).

Keller, W. (2004). International technology diffusion. *Journal of Economic Literature*, 42(3):752–782.

Kuusi, T., Kulvik, M., and Junnonen, J.-M. (2022). Productivity growth in construction value chains. *International Productivity Monitor*, 42:3–32.

Lee, E. and Yi, K.-M. (2018). Global value chains and inequality with endogenous labor supply. *Journal of International Economics*, 115:223–241.

Leontief, W. (1936). Quantitative input and output relations in the economic system of the united states. *Review of Economic Statistics*, 18(3):105–125.

Los, B., Timmer, M. P., and de Vries, G. J. (2016). Tracing value-added and double counting in gross exports: Comment. *American Economic Review*, 106(7):1958–1966.

Mammen, E. (1993). Bootstrap and wild bootstrap for high dimensional linear models. *Annals of Statistics*, 21(1):255–285.

Novy, D. (2013). Gravity redux: Measuring international trade costs with panel data. *Economic Inquiry*, 51(1):101–121.

Prescott, E. C. and Visscher, M. (1980). Organization capital. *Journal of Political Economy*, 88(3):446–461.

Reijnders, L. S. M. and de Vries, G. J. (2018). Technology, offshoring and the rise of non-routine jobs. *Journal of Development Economics*, 135:412–432.

Sant'Anna, P. H. C. and Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1):101–122.

Timmer, M. (2017). Productivity measurement in global value chains. *International Productivity Monitor*, 33:182–193.

Timmer, M. P. and Ye, X. (2020). Accounting for growth and productivity in global value chains. In *Measuring Economic Growth and Productivity*, pages 413–426. Elsevier.

van Ark, B., de Vries, K., and Erumban, A. (2024). Are intangibles running out of steam? *International Productivity Monitor*, 46:38–59.

Wolff, E. N. (1994). Productivity measurement within an input-output framework. *Regional Science and Urban Economics*, 24(1):75–92.

Zeng, K. and Lu, Y. (2016). Variation in bilateral investment treaty provisions and foreign direct investment flows to china, 1997–2011. *International Interactions*, 42(5):820–848.

A Appendix

A.1 VC TFP measurement: additional analysis

Measurement error relative to conventional TFP. To further assess the robustness of the VC TFP measure, we compare annual VC TFP growth with conventional EU KLEMS TFP growth at the same country–industry–year level. We define the measurement wedge as

$$\varepsilon_{it} = g_{it}^{VC} - g_{it}^{KLEMS}, \tag{19}$$

and its squared value as

$$\varepsilon_{it}^2 = \left(g_{it}^{VC} - g_{it}^{KLEMS}\right)^2. \tag{20}$$

The purpose of this exercise is not to treat KLEMS TFP as a "true" benchmark. Rather, it is a validation exercise that asks two questions. First, is the wedge between the two measures systematically larger in value chains with a larger role for market services, where standard industry-based TFP measures are more likely to miss reorganization and upstream service contributions? Second, does the wedge—or its volatility—increase mechanically with the foreign share of the value chain, which would be a concern given that the VC TFP measure uses foreign input data for the upstream part of the chain?

Table 9 reports fixed-effects regressions for the signed wedge ε_{it}. The results point to three broad findings. First, the wedge is smaller when the final producer accounts for a larger share of value added: the coefficient on the final producer share is negative in the full sample. This means that the difference between VC TFP and conventional TFP is largest in chains where a greater share of value added originates outside the final producer industry. Second, the wedge is systematically larger in value chains with a larger market-services component. The coefficient on the domestic market-services share is positive and statistically significant in the full sample, in 1995–2007, and in 2008–2017. By contrast, manufacturing shares are small and imprecisely estimated. This pattern is consistent with the broader interpretation of the paper: conventional industry-level

TFP measures are particularly restrictive in service-intensive and networked value chains, where productivity improvements operate through logistics, organization, finance, marketing, and other market services that are not confined to the final producer industry. This interpretation is also in line with the value-added composition analysis (Appendix Tables 17 and 19), which show a large and increasing role for market services in European value chains.

Table 10 turns to the squared wedge ε_{it}^2, which captures volatility in the difference between VC TFP and conventional TFP. Here the picture is notably different. The market-services effect that is visible in the signed wedge is not statistically discernible in the squared wedge, implying that service intensity is associated primarily with a *systematic level difference* between VC TFP and conventional TFP, not with unstable measurement. Most importantly for the use of foreign upstream information, there is no systematic evidence that a larger foreign component of the value chain produces either excess volatility or mechanical bias in the new measure. In the ε_{it} regressions, the coefficients on foreign manufacturing, foreign market services, and foreign other services are statistically insignificant throughout. In the ε_{it}^2 regressions, neither the overall foreign shares nor the foreign service share are statistically significant. Thus, while the level difference between VC TFP and conventional TFP is larger in service-intensive chains, the additional reliance on foreign upstream data does not appear to create a mechanically more volatile or systematically biased measure. If anything, the results suggest that the main difference between the two productivity concepts reflects the economic content of value chains rather than noise coming from the foreign part of the data.

Direct comparison with earlier VC TFP exercises. The regression evidence above is complemented by direct comparisons to earlier applications of value-chain productivity accounting. We revisit the German automotive industry discussed by Timmer (2017) and the construction example in Kuusi et al. (2022). These examples are useful because they provide transparent benchmarks where a value-chain perspective should matter most: automotive production relies heavily on inter-

Table 9: VC TFP measurement wedge relative to conventional TFP and value-added composition

	Dependent variable: $\varepsilon = g^{VC} - g^{KLEMS}$			
	Own share 1995–2017	Composition 1995–2017	Composition 1995–2007	Composition 2008–2017
Final producer share	-0.15**			
	(0.05)			
Domestic manufacturing share		-0.02	-0.40	0.60
		(0.13)	(0.29)	(0.38)
Foreign manufacturing share		0.00	0.45	-0.18
		(0.19)	(0.35)	(0.62)
Domestic market services share		0.24*	0.56*	0.41**
		(0.10)	(0.26)	(0.15)
Foreign market services share		0.12	0.36	0.15
		(0.27)	(0.60)	(0.59)
Domestic other services share		0.39	0.67	0.19
		(0.28)	(0.69)	(0.53)
Foreign other services share		0.24	-16.35	6.28
		(8.09)	(16.69)	(15.18)
Year fixed effects	Yes	Yes	Yes	Yes
Value-chain fixed effects	Yes	Yes	Yes	Yes
Observations	7,393	7,393	3,581	3,812
Groups	392	392	336	392
Within R^2	0.010	0.010	0.019	0.013

Notes: Fixed-effects estimates with robust standard errors clustered at the value-chain panel identifier (`id`). The dependent variable ε is the annual difference between VC TFP growth and EU KLEMS TFP growth. "Final producer share" is the value-added share of the final producer industry in the chain. The composition variables are domestic and foreign value-added shares of upstream manufacturing, market services, and other services. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Volatility of the VC TFP measurement wedge and value-added composition

	Dependent variable: $\varepsilon^2 = \left(g^{VC} - g^{KLEMS}\right)^2$	
	Own share 1995–2017	Composition 1995–2017
Final producer share	-0.14	
	(0.09)	
Domestic manufacturing share		0.17
		(0.10)
Foreign manufacturing share		0.06
		(0.13)
Domestic market services share		0.06
		(0.10)
Foreign market services share		0.16
		(0.20)
Domestic other services share		0.43
		(0.31)
Foreign other services share		-1.20
		(4.94)
Year fixed effects	Yes	Yes
Value-chain fixed effects	Yes	Yes
Observations	7,393	7,393
Groups	392	392
Within R^2	0.006	0.005

Notes: Fixed-effects estimates with robust standard errors clustered at the value-chain panel identifier (`id`). The dependent variable ε^2 is the square of the annual difference between VC TFP growth and EU KLEMS TFP growth. "Final producer share" is the value-added share of the final producer industry in the chain. The composition variables are domestic and foreign value-added shares of upstream manufacturing, market services, and other services. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

nationally sourced components and business services, while construction productivity increasingly depends on upstream factory-based production and prefabrication rather than on-site activity.

Overall, the comparisons increase confidence that the VC TFP series capture economically meaningful patterns rather than idiosyncratic noise. In the German automotive industry in 1995–2007, we confirm the main message in Timmer (2017): a substantial part of productivity growth originates outside the final industry. In our data, output grew by 6.0% annually, input contributions account for 3.7 percentage points, and VC TFP growth amounts to 2.4 percentage points. Of this, only 0.8 percentage points come from the final producer industry itself, while 1.5 percentage points originate from the rest of the value chain. Relative to Timmer (2017), the input contribution is very similar, but the share of final-industry TFP is smaller in our data, suggesting a larger role for foreign upstream contributions.

The same exercise also shows how the value-chain perspective helps interpret the post-2008 slowdown. In German automotive manufacturing, output growth slows to 2.3% in 2008–2017, input growth falls markedly, and upstream VC TFP weakens from 1.5 to 0.3 percentage points. At the European level, automotive output is essentially stagnant in 2008–2017, but VC TFP remains positive because of productivity gains in the broader chain. Finally, in construction, our findings line up with Kuusi et al. (2022): VC TFP provides a less pessimistic picture than conventional industry-based productivity measures because upstream productivity contributions are positive even when the final industry's own TFP contribution is negative. This is exactly the type of environment where value-chain accounting is informative, since technological progress often shifts activity from the final producer to upstream suppliers and prefabricated production stages.

Heterogeneity in VC productivity dynamics across countries. The aggregate slowdown masks substantial cross-country heterogeneity. Some countries maintained positive VC TFP growth after 2008 by reducing inputs while preserving productivity, while others experienced weak VC TFP performance alongside weak growth. The tables below provide country-level decompositions

Table 11: Corollaries with earlier VC TFP studies and extensions

	Output growth	Input contr.	VC TFP	Final TFP	Other VC TFP
German automot., 1995–2007	6.0%	3.7	2.4	0.8	1.5
German automot., 2008–2017	2.3%	1.0	1.3	0.9	0.3
European automot., 2008–2017	-0.1%	-1.3	1.2	0.4	0.8
European construction, 1995–2017	0.5%	0.6	-0.1	-0.3	0.2

Notes: Entries report average annual growth rates (percentage points). "Final TFP" denotes EU KLEMS TFP growth in the final producer industry and "Other VC TFP" the contribution of the remaining value chain. The automotive example revisits Timmer (2017); the construction example revisits Kuusi et al. (2022) using the present dataset.

for market services and manufacturing. The discussion emphasizes broad patterns rather than country-by-country narratives.

Table 12: Output growth and its decomposition for market services by country (rows).

	1995-2007			2008-2017		
	Output growth	Input contr.	VC TFP	Output growth	Input contr.	VC TFP
AUT	5.3 %	3.9	1.5	0.6 %	1.0	-0.4
BEL	2.6 %	2.9	-0.3	1.6 %	2.6	-0.9
DEU	4.3 %	3.3	1.1	2.1 %	1.6	0.5
DNK	5.3 %	4.7	0.6	1.3 %	0.7	0.6
ESP	5.2 %	5.5	-0.3	1.0 %	1.2	-0.1
FIN	4.8 %	2.8	2.0	1.0 %	1.1	0.0
FRA	4.4 %	3.7	0.7	1.7 %	2.2	-0.5
GBR	4.0 %	2.3	1.7	1.5 %	1.9	-0.4
GRC				-1.9 %	1.4	-3.3
ITA	3.4 %	3.8	-0.4	-0.3 %	0.3	-0.6
LUX				3.7 %	3.8	-0.1
NLD	5.1 %	4.0	1.1	1.9 %	2.7	-0.8
SWE	4.5 %	4.0	0.5	2.1 %	1.7	0.4
Total	**2.8 %**	**2.6**	**0.2**	**1.1 %**	**1.6**	**-0.5**

Notes: Country codes are ISO-3. Entries are average annual growth rates (percentage points). Output growth is decomposed into the contribution of aggregate primary inputs used across the full value chain and VC TFP. A tilde (~) denotes unavailable data for the country-sector in the sample window.

Two broad patterns stand out. First, post-2008 VC TFP performance is often associated with input contraction, consistent with restructuring and reallocation. Second, a subset of countries maintains positive VC TFP growth in both market services and manufacturing, while others exhibit persistently weak VC TFP growth across sectors. This motivates examining the composition of inputs and the distribution of productivity across value chain segments.

Table 13: Output growth and its decomposition for manufacturing by country.

	1995-2007			2008-2017		
	Output growth	Input contr.	VC TFP	Output growth	Input contr.	VC TFP
AUT	4.0 %	2.3	1.8	1.2 %	1.1	0.1
BEL	1.2 %	0.5	0.7	-0.9 %	-0.1	-0.8
DEU	2.7 %	1.1	1.7	0.4 %	0.0	0.3
DNK	1.2 %	0.8	0.3	-0.8 %	-1.3	0.5
ESP	3.1 %	2.5	0.5	-0.3 %	-2.3	2.0
FIN	4.0 %	2.2	1.8	-0.8 %	-1.2	0.4
FRA	2.2 %	1.3	0.9	-0.8 %	-0.5	-0.2
GBR	1.4 %	-1.5	2.9	-0.8 %	-0.8	0.0
GRC				-4.0 %	-3.8	-0.2
ITA	2.5 %	2.5	0.0	-1.4 %	-1.3	-0.1
LUX				-1.9 %	-1.4	-0.5
NLD	2.4 %	1.1	1.3	0.8 %	0.9	-0.1
Average	**2.5 %**	**1.3**	**1.2**	**-0.7 %**	**-0.8**	**0.1**

Notes: Country codes are ISO-3. Entries are average annual growth rates (percentage points). Output growth is decomposed into the contribution of aggregate primary inputs used across the full value chain and VC TFP. A tilde (~) denotes unavailable data for the country-sector in the sample window.

Final producer vs. upstream heterogeneity after 2008. Country-level heterogeneity is also visible in the decomposition of VC TFP into final producer vs upstream components in 2008–2017. Table 14 summarizes these patterns. In some countries, final producer TFP remains positive while upstream TFP is negative, consistent with weak diffusion of productivity improvements along the chain. In others, both components are negative, consistent with a broader productivity malaise.

Input composition: domestic vs. foreign labor and capital. To interpret why input growth and VC TFP changed differently across sectors, we decompose input contributions into domestic and foreign labor and capital components. Table 15 shows a broad decline in input growth after 2008, especially in manufacturing. In market services, domestic capital contributions remain comparatively stable, which is consistent with the view that much of post-2008 growth effort was capital deepening in services, even as value chain TFP weakened.

Table 16 summarizes domestic vs foreign input contributions by country. The foreign component varies substantially across countries, consistent with differences in exposure to GVCs and the role of international services. These patterns provide context for why the BIT shock and the intangible spillovers matter differently across countries.

Table 14: TFP decomposition to final industry TFP and the rest of the value chain TFP by country (rows) in 2008-2017.

	Market services 08-		Manufacturing 08-	
	Final TFP	Other VC TFP	Final TFP	Other VC TFP
AUT	0.1	-0.5	0.0	0.1
BEL	-0.2	-0.7	0.3	-1.1
DEU	0.4	0.1	0.3	0.0
DNK	0.5	0.1	0.8	-0.3
ESP	-0.3	0.2	0.3	1.7
FIN	0.2	-0.3	0.3	0.1
FRA	0.0	-0.5	0.3	-0.5
GBR	1.3	-1.8	0.8	-0.8
GRC	-1.6	-1.7	-0.6	0.3
ITA	-0.3	-0.3	0.0	0.0
LUX	0.1	-0.2	2.9	-3.4
NLD	0.0	-0.8	0.3	-0.4
SWE	0.2	0.1		

Notes: Country codes are ISO-3. Entries decompose VC TFP into final-producer TFP and upstream VC TFP (percentage points per year) for 2008–2017.

Table 15: Decomposition of input contributions to output growth by sector and period (rows), percentage points

		Labor		Capital	
	Total	Domestic	Foreign	Domestic	Foreign
Manufacturing					
95-07	1.3	-0.1	0.2	0.7	0.6
08-17	-0.8	-0.6	-0.3	0.1	0.1
Market services					
95-07	3.8	1.2	0.3	1.9	0.4
08-17	1.6	0.3	0.1	0.9	0.3
Other services					
95-07	2.2	1.2	0.1	0.7	0.2
08-17	1.5	0.8	0.0	0.5	0.1
Utilities + cons					
95-07	2.5	0.5	0.3	1.2	0.5
08-17	0.8	0.0	-0.2	0.9	0.1

Notes: Entries report average annual changes (percentage points) in the contributions of labor and capital inputs used across the value chain. “Change” refers to the difference between 2008–2017 and 1995–2007 averages.

Table 16: Domestic and foreign input contribution to output growth by country (row), percentage points

	Manufacturing		Market services	
	Domestic	Foreign	Domestic	Foreign
AUT	0.3	0.8	0.5	0.5
BEL	-0.2	0.1	1.6	0.9
DEU	0.0	0.0	1.1	0.5
DNK	-0.8	-0.5	0.6	0.1
ESP	-1.5	-0.8	1.4	-0.2
FIN	-0.5	-0.8	0.8	0.3
FRA	-0.4	-0.2	1.9	0.3
GBR	-0.4	-0.5	1.8	0.1
GRC	-2.5	-1.2	1.6	-0.2
ITA	-0.7	-0.6	0.5	-0.2
LUX	-0.1	-1.4	1.3	2.5
NLD	0.1	0.7	1.7	0.9
SWE			1.6	0.1

Notes: Entries report average annual contributions (percentage points) of domestic vs. foreign labor and capital inputs to output growth by country. Country codes are ISO-3.

Value-added composition in the value-chain data. We also examine the value-added composition in our value-chain data. The results can be found in Table 17. This analysis focuses on manufacturing, market services, other services, and the construction and utilities sector. The value-added breakdown includes contributions from the final producer industry and all upstream industries.[15]

Between 1995 and 2017, the final producer industry accounted for about 56.5% of the value-added in the value chain, with the majority originating from the domestic part of the chain. Market services contributed 25.7% of the value-added, manufacturing 8.4%, and final production 4.5%. The data reveals consistent organization across sectors, with the final producer industry's share being the largest, varying between 40% and 76%. The highest shares were found in other services,

[15] Manufacturing includes all manufacturing activities (ISIC class C). Utilities and Construction encompasses utilities (such as electricity, gas, steam, and air conditioning supply), water supply, sewerage, waste management, and remediation activities, as well as construction (ISIC class D-F). Market Services include wholesale and retail trade, transportation and storage, accommodation and food service activities, information and communication, financial and insurance activities, real estate activities, professional, scientific and technical activities, and administrative and support service activities (ISIC class G-N). Other Services cover public administration and defense, education, human health and social work activities, arts, entertainment and recreation, other service activities, activities of households as employers, and activities of extraterritorial organizations and bodies (ISIC class O-U). We also include agriculture, forestry, fishing, and mining activities (ISIC class A, B) as provider of inputs

while the lowest were in manufacturing. The role of market services in value added also remains consistently high, with the highest share (outside the final producer industry) in the manufacturing sector.

From 1995 to 2017, the most significant change has been the increase in the private service sector's value-added share, which rose by 4.2 percentage points. Meanwhile, the shares of the final industry, manufacturing, and the final sector decreased by roughly the same amount. This change predominantly occurred in the foreign part of the value chain, as shown in Table 20. Professional, scientific, technical, administrative, support service activities, financial, insurance, and trade activities contributed significantly to the increase in the services sector.

Table 17: Value-added composition for manufacturing, market services, other services, and construction and utilities sector (columns), as well as unweighted averages

	Average	Manufacturing	Market services	Other services	Util. & construction
Final industry	56.5 %	40.4 %	57.7 %	76.1 %	52.1 %
Manufacturing	8.4 %	12.7 %	6.4 %	4.2 %	10.4 %
Market services	25.7 %	32.7 %	28.6 %	15.1 %	26.5 %
Other services	1.8 %	1.7 %	2.0 %	1.3 %	2.1 %
Final	4.5 %	8.1 %	2.5 %	1.2 %	6.2 %
Util. & construct.	3.0 %	4.1 %	2.7 %	2.2 %	2.8 %

Notes: Value-added (VA) shares are computed from OECD ICIO and reported as percentages of gross output. "Final producer" denotes VA generated in the final-producer industry; "Other industries" aggregates VA generated upstream.

Table 18: Value-added composition changes 1995-2017 for manufacturing, market services, other services, and construction and utilities sector (columns), as well as unweighted averages, percentage points

	Average	Manufacturing	Market services	Other services	Util. & construction
Final industry	-1.6	-3.7	-1.2	-1.4	-0.2
Manufacturing	-1.1	-0.1	-1.4	-0.6	-2.4
Market services	4.2	4.7	3.7	2.5	6.0
Other services	0.5	0.7	0.4	0.6	0.4
Final	-2.1	-2.1	-0.9	-0.7	-4.6
Util. & construct.	0.0	0.5	-0.7	-0.4	0.8

Notes: Entries report changes (percentage points) in value-added shares between 1995 and 2017. "Final producer" denotes VA generated in the final-producer industry; "Other industries" aggregates VA generated upstream.

Table 19: Foreign and domestic value-added decomposition for manufacturing, market services, other services, and construction and utilities sector (columns)

		Average	Manufacturing	Market services	Other services	Util. & construction
Final industry	dom	55.3 %	37.5 %	55.8 %	76.0 %	51.7 %
	for	1.3 %	2.9 %	1.9 %	0.0 %	0.4 %
Manufacturing	dom	4.3 %	5.8 %	3.4 %	2.2 %	5.8 %
	for	4.1 %	6.8 %	3.0 %	2.0 %	4.6 %
Market services	dom	15.8 %	17.6 %	18.8 %	10.4 %	16.5 %
	for	9.9 %	15.0 %	9.8 %	4.7 %	10.0 %
Other services	dom	1.4 %	1.1 %	1.6 %	1.1 %	1.7 %
	for	0.4 %	0.6 %	0.4 %	0.1 %	0.4 %
Final	dom	1.1 %	2.3 %	0.5 %	0.3 %	1.3 %
	for	3.4 %	5.9 %	2.0 %	0.8 %	4.9 %
Util. & construct.	dom	2.3 %	2.8 %	2.2 %	2.0 %	2.3 %
	for	0.6 %	1.3 %	0.5 %	0.3 %	0.5 %
Total	dom	80.2 %	67.2 %	82.3 %	92.0 %	79.3 %
	for	19.7 %	32.4 %	17.7 %	8.0 %	20.7 %

Notes: Decomposes value-added into domestic vs. foreign components (shares of gross output, percent) using OECD ICIO. “Final producer” denotes the final-producer industry; “Other industries” aggregates upstream sectors.

Table 20: Foreign and domestic value-added composition changes 1995-2017 for manufacturing, market services, other services, and construction and utilities sector (columns), percentage points

		Average	Manufacturing	Market services	Other services	Util. & construction
Final industry	dom	-2.0	-3.6	-2.7	-1.4	-0.1
	for	0.4	0.0	1.6	0.0	-0.1
Manufacturing	dom	-1.0	-0.6	-1.1	-0.4	-2.0
	for	-0.1	0.4	-0.2	-0.2	-0.4
Market services	dom	-0.5	-1.3	-1.4	0.0	0.8
	for	4.7	6.0	5.1	2.6	5.2
Other services	dom	0.3	0.5	0.1	0.5	0.3
	for	0.2	0.3	0.3	0.1	0.2
Final	dom	-0.4	-0.3	-0.2	-0.2	-0.8
	for	-1.7	-1.9	-0.7	-0.5	-3.8
Util. & construct.	dom	0.0	0.4	-0.7	-0.4	0.7
	for	0.1	0.1	0.1	0.0	0.0
Total	dom	-3.5	-4.9	-6.1	-2.0	-1.0

Notes: Entries report changes (percentage points) in domestic and foreign value-added shares between 1995 and 2017 using OECD ICIO. “Final producer” denotes the final-producer industry; “Other industries” aggregates upstream sectors.

A.2 Event study: robustness, interpretation, and heterogeneity

Table 21: Pre-treatment parallel trends: joint test and event-study estimates

	VC TFP		Bus. intangibles		Tech. intangibles	
Horizon	ATT	SE	ATT	SE	ATT	SE
Pre-treatment periods						
$k = -7$	0.004	(0.004)	−0.000	(0.009)	−0.049	(0.031)
$k = -6$	0.006	(0.006)	0.006	(0.008)	0.018	(0.013)
$k = -5$	−0.004	(0.007)	0.004	(0.010)	0.041	(0.019)
$k = -4$	0.001	(0.005)	−0.003	(0.011)	−0.025	(0.019)
$k = -3$	−0.001	(0.005)	−0.015	(0.008)	0.005	(0.020)
$k = -2$	0.003	(0.006)	−0.014	(0.014)	−0.046	(0.014)
Joint pre-trend test (H_0: all pre-treatment ATTs = 0)						
χ^2 (df=6)	8.220		12.073		25.040	
p-value	0.222		0.060		0.000	

Notes: Callaway–Sant'Anna (2021) CS-DiD DRIPW estimator. Mammen wild bootstrap standard errors (999 reps), resampled at the country–industry panel-unit level; treatment timing varies across 15 countries.
Joint test: Wald chi-squared test that all six pre-treatment ATTs are jointly zero. $k = -1$ is the omitted baseline.
Failure to reject the joint null for VC TFP ($p = 0.22$) supports the parallel-trends assumption for the main TFP outcome. Business intangibles borderline ($p = 0.06$); Tech intangibles rejects ($p < 0.001$), consistent with the insignificant $t + 7$ ATT reported for that outcome.

Additional BIT interpretation and heterogeneity. The BIT evidence in Section 5 shows a positive VC TFP response and a clear increase in business-related intangibles, but the magnitude and location of effects vary across sectors and value chain segments. Two interpretive points are useful for the broader narrative.

First, the VC TFP response is more pronounced in market services than in manufacturing when comparing VC TFP to traditional TFP (Figures 1–5). This is consistent with the idea that service-intensive value chains rely heavily on reorganizing tasks and relationships (logistics, finance, management, and other market services), which may be poorly captured in traditional productivity measures and more sensitive to intangible capability. Second, the decomposition into final producer vs upstream TFP indicates that market services exhibit clearer upstream improvements after BITs than manufacturing (Figure 2), which is consistent with services being intermediated through broader network linkages.

Figure 5 shows that the difference between VC TFP and traditional TFP is more pronounced in market services, consistent with the idea that value chain linkages and input mismeasurement are more salient in service-intensive chains.

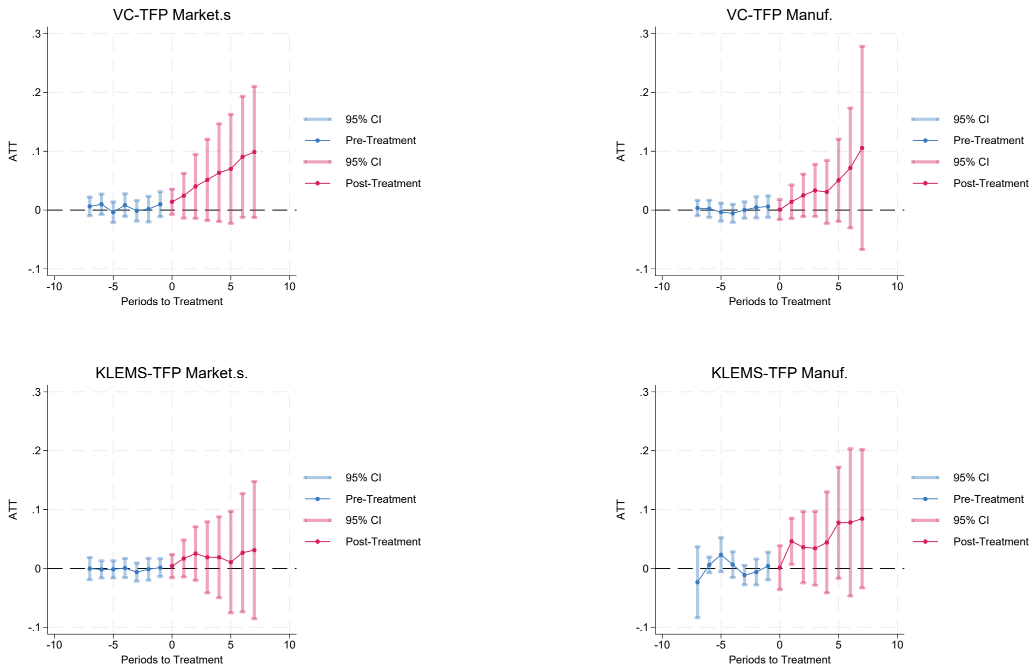

Figure 5: Development of TFP after BITs agreement enters into force: manufacturing vs. market services

Note: The dots are event coefficients and the bars their bootstrapped 95 % confidence intervals clustered at unit level. The pre-treatment coefficients are blue and the post-treatment coefficients red.

Robustness of the event study analysis. The treatment-effect heterogeneity is limited. Examining manufacturing and market services separately yields broadly similar point estimates to the pooled baseline. Pre-event coefficients are statistically insignificant in almost all subgroups, and post-event coefficients show a consistent positive pattern across groups, though individual group estimates are imprecise.

Identification rests on treatment timing in 15 countries, which is at the lower bound of what is typically required for cluster-robust inference at the level of policy assignment. Re-doing inference with clustering at the country (policy) level rather than at the country–industry panel unit renders the post-event coefficients insignificant at the 5% level, though the confidence intervals remain predominantly positive; we regard this as the honest characterization of the precision available at the policy level. Alternative estimators, including Borusyak et al. (2024) and two-way fixed effects, produce estimates close to the baseline.

Extending the event window to 13 years before and after treatment, post-event coefficients average 0.13 for years 9–13, broadly in line with the seven-year estimates. Pre-event coefficients remain close to zero throughout the pre-period, providing no evidence of anticipation effects.

To assess whether results are driven by a single influential country, we estimate the model leaving one country out at a time. The point estimates and pre-trend patterns are stable across these leave-one-out samples, indicating that no single country accounts for the main findings.

As an additional robustness check, we restrict the sample to countries that eventually signed a BIT with China. This addresses the concern that our baseline estimates could partly reflect unobserved and persistent differences between countries that signed a BIT and those that never did, rather than the effects of the treaty itself. In the treated-only sample, identification comes from differences in treaty timing across countries that all eventually receive the treatment, so the comparison is made within a more homogeneous set of economies and is less exposed to selection based on permanent cross-country characteristics.

Results are reported in Figure 6. We continue to find positive and statistically significant productivity effects of broadly similar magnitude to those in the baseline specification. This suggests that the main findings are not driven by a simple distinction between BIT-signing and non-signing countries. At the same time, this exercise does not by itself rule out all sources of endogeneity in treaty timing, since the decision of when to ratify a BIT may still be related to time-varying country-specific factors. Even so, the similarity of the treated-only estimates to the baseline results provides additional support for the interpretation that BITs are associated with economically meaningful productivity gains.

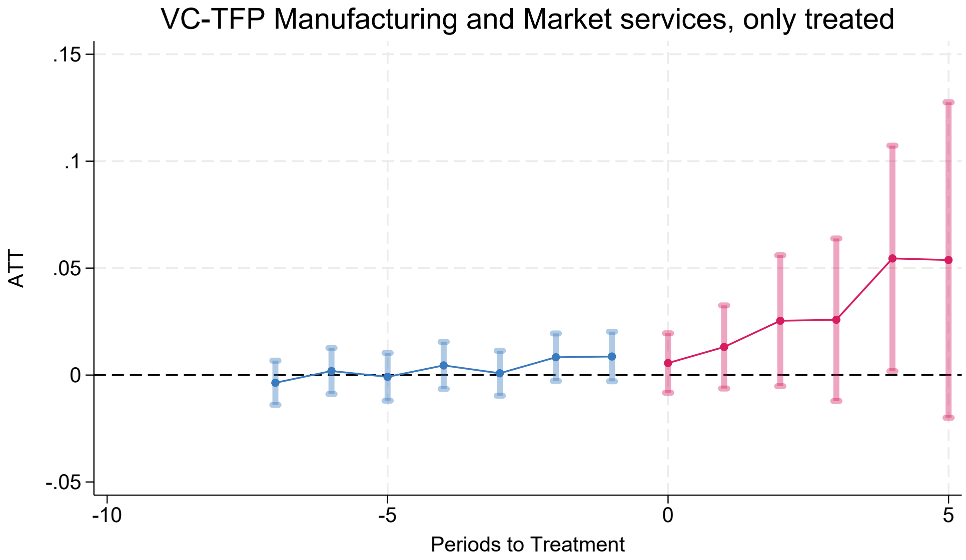

Figure 6: Development of TFP after BITs agreement enters into force: countries that eventually signed a BIT with China

Note: The dots are event coefficients and the bars their bootstrapped 95 % confidence intervals clustered at unit level. The pre-treatment coefficients are blue and the post-treatment coefficients red.

A.3 Treatment intensity: heterogeneity by China exposure

The country-level CS-DiD in Section 5 exploits variation in BIT timing across 15 countries (9 treated, 6 never-treated). This appendix pushes the identifying variation to the industry level by asking whether industries that were more integrated with Chinese value chains before the BITs entered into force responded more strongly to the treaty than less-exposed industries.

Intensity variable. The intensity weight is each industry's pre-treaty average share of Chinese value added in its total output (TIVA share):

$$\bar{s}_i^{\mathrm{CHN}} = \frac{1}{C}\sum_c \frac{1}{9}\sum_{t=1995}^{2003} s_{ict}^{\mathrm{CHN}}, \tag{21}$$

where s_{ict}^{CHN} is China's value-added share in industry i, country c, year t. We average over all final-producer countries c. Being constructed entirely before any BIT entered into force, $\bar{s}_i^{\text{CHN}}$ is predetermined with respect to the treatment.

Estimation strategy. We partition the sample at the industry-level median of $\bar{s}_i^{\text{CHN}}$ into a "high-exposure" subsample (industries strictly above the median) and a "low-exposure" subsample (industries at or below the median), and estimate the Callaway–Sant'Anna group-time ATT estimator with doubly-robust inverse-propensity weighting (Callaway and Sant'Anna, 2021; Sant'Anna and Zhao, 2020) separately in each partition. The panel unit is the country–industry cell; time is annual; treatment cohort "g" is the year of BIT entry into force for the country (never-treated for control countries). We use the not-yet-treated comparison group and set the pre-treatment reference period to $k = -1$.

Each partition-specific CS-DiD is identified under the same assumptions as the pooled baseline: no anticipation of the BIT, and parallel trends of the untreated potential outcome against the not-yet-treated comparison group at every event time. We report both the long-run ATT at k=+7 and a joint Wald test on the six pre-treatment ATTs ($\tau_{-7}, \ldots, \tau_{-2}$).

Inference. Standard errors are computed via the Mammen wild bootstrap (Mammen, 1993) with 999 replications, clustered at the country–industry panel-unit level, which corresponds to resampling the influence-function contributions at the unit level. The seed is fixed for reproducibility. The high-exposure partition contains 213 unique country–industry clusters (4,461 observations across 22 years); the low-exposure partition contains 218 clusters (4,545 observations). Both partitions retain the full staggered treatment structure across the nine treated countries (the Netherlands, Sweden, Germany, Finland, Portugal, Spain, Belgium, Luxembourg, France; treated years 2004–2010) and the six never-treated control countries. A direct high-versus-low difference test treats the two partition estimates as independent, on the grounds that the partitions are drawn from disjoint industry subsets, and combines them via a Wald-type z-statistic $z = \left(\widehat{\text{ATT}}^{\text{HI}}(+7) - \widehat{\text{ATT}}^{\text{LO}}(+7)\right) / \sqrt{\widehat{\text{SE}}_{\text{HI}}^{2} + \widehat{\text{SE}}_{\text{LO}}^{2}}$. The independence assumption requires a qualification. Treatment is assigned at the country level, so both partitions within a country share the same BIT timing and are exposed to the same country-level shocks; the two partition estimates are therefore positively correlated, and ignoring that covariance understates the standard error of the difference. The direction of the bias works against the conclusion we draw: a correctly signed covariance term would make the difference test *less* likely to reject, and the test already fails to reject. The substantive finding of no detectable heterogeneity along the China-exposure margin is therefore unaffected, though the reported z-statistics should be read as approximate.

Table 22: Partition-based CS-DiD: heterogeneity by China exposure

	High China exposure		Low China exposure	
Outcome	ATT($k=+7$) (SE)	Joint pre-trend $\chi^2(6)$ [p]	ATT($k=+7$) (SE)	Joint pre-trend $\chi^2(6)$ [p]
Cum. VC TFP	0.097*	8.98	0.108***	16.68
	(0.058)	[0.175]	(0.038)	[0.011]
Cum. KLEMS TFP	0.029	21.41	0.091*	5.05
	(0.028)	[0.002]	(0.052)	[0.538]
Log total intang. stock	0.072*	7.60	0.037	11.76
	(0.038)	[0.269]	(0.047)	[0.068]
Log bus. intang. stock	0.067*	25.11	0.079*	6.47
	(0.040)	[0.000]	(0.042)	[0.372]
HIGH − LOW difference (independent-samples Wald test):				
Cum. VC TFP	$\Delta=-0.011$, SE=0.070, $z=-0.16$, p=0.870			
Cum. KLEMS TFP	$\Delta=-0.063$, SE=0.059, $z=-1.05$, p=0.293			
Log total intang. stock	Δ=0.035, SE=0.060, z=0.58, p=0.563			
Log bus. intang. stock	$\Delta=-0.012$, SE=0.058, $z=-0.21$, p=0.836			

Notes: Callaway–Sant'Anna (2021) DRIPW estimator; not-yet-treated comparison group; wild (Mammen) bootstrap standard errors with 999 replications, clustered at the country–industry panel-unit level. Sample split at the industry-level median of $\bar{s}_i^{\mathrm{CHN}}$: "high"/"low" contain the industries strictly above/at-or-below the median. The high-exposure partition has 213 country–industry clusters (4,461 observations); the low-exposure partition has 218 clusters (4,545 observations). ATT(k=+7) is the long-run treatment effect; standard errors in parentheses. "Joint χ^2" tests $H_0 : \tau_{-7} = \cdots = \tau_{-2} = 0$ within the partition (df = 6), with p-value in brackets. The difference row at the bottom computes a Wald-type z-statistic on $\widehat{\mathrm{ATT}}^{\mathrm{HI}}(+7) - \widehat{\mathrm{ATT}}^{\mathrm{LO}}(+7)$ under independence of the two subsample estimates. Significance stars are attached to the ATT and Δ point estimates from the two-sided Wald p-value: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Partition-based results. Table 22 reports the partition-based estimates for the four outcomes. Three points stand out. First, all long-run ATTs are positive in both partitions: the BIT is associated with cumulative productivity gains and higher intangible stocks in both high- and low-exposure industries. Second, the high-versus-low differences are small in magnitude and never statistically significant at conventional levels: the largest differential is −6.3 log points for cumulative KLEMS TFP ($p = 0.29$), and the VC TFP and business-intangible differentials are close to zero ($|\Delta| \approx 0.01$). Third, the pre-trend picture is heterogeneous across partitions: the high-exposure subsample shows a substantive pre-trend violation for cumulative KLEMS TFP ($\chi^2(6)$=21.4, p=0.002) and for the business-intangible stock ($\chi^2(6)$=25.1, $p < 0.001$), while the low-exposure subsample fails to reject parallel trends for both intangible outcomes.

Taken together, once identification is confined to within-partition comparisons that respect the CS-DiD building-block structure, there is no evidence of statistically distinguishable heterogeneity in the BIT effect along the China-exposure dimension. The pre-trend violations concentrated in the high-exposure subsample suggest that industries with strong pre-BIT Chinese value-chain integration were already on a distinctive trajectory, which should temper the causal interpretation of the high-exposure post-BIT ATTs even though the point estimates remain positive.

Clustering-level robustness. Splitting the sample does not create additional country-level policy experiments: counting country–partition cells as 30 clusters would not remove within-country dependence, since BIT timing and country-level shocks are common to both partitions of a given country. What the partition design does allow is a check on the level of clustering. Because the pooled event study resamples wild-bootstrap replications at the country–industry panel-unit level, we can re-estimate the same objects under coarser (country-level) clustering, which matches the level at which BIT treatment is actually assigned, and see how much precision is attributable to the finer unit.

Table 23 reports two additional specifications: (i) a *country-clustered per-partition* exercise in which the industry-country panel is collapsed to country-year cell means within each partition and the CS-DiD is re-estimated with the country as the panel unit (15 country clusters per partition); and (ii) a *stacked* specification in which the collapsed data are pooled across both partitions and the panel unit is set to the country–partition cell, giving 30 unit-panel clusters. As expected, the country-clustered ATTs are less precisely estimated than the baseline industry-cluster version. Nonetheless, the two more informative outcomes—cumulative VC TFP and log business intangibles—retain positive point estimates and marginal significance in the LO and stacked specifications, and the qualitative conclusion is unchanged: none of the HI − LO differences is statistically significant at conventional levels ($|z| < 1$ in every case).

Table 23: Cluster-level robustness for the partition CS-DiD

	Country-cluster (15)				Stacked (30 c-p)
Outcome	HI ATT (SE)	LO ATT (SE)	Δ	z [p]	Pooled ATT (SE)
Cum. VC TFP	0.109	0.118**	-0.009	-0.09 [0.929]	0.114**
	(0.086)	(0.052)			(0.050)
Cum. KLEMS TFP	0.033	0.086	-0.054	-0.69 [0.490]	0.059
	(0.055)	(0.055)			(0.039)
Log total intang. stock	0.014	0.109*	-0.096	-0.84 [0.403]	0.061
	(0.095)	(0.064)			(0.059)
Log bus. intang. stock	0.085	0.100	-0.014	-0.14 [0.891]	0.093*
	(0.082)	(0.063)			(0.052)

Notes: All columns use the Callaway–Sant'Anna DRIPW estimator with the not-yet-treated comparison group, Mammen wild bootstrap standard errors (999 reps), and event window $k \in [-7, +7]$. “Country-cluster (15)”: the industry-country panel is collapsed to country-year cell means within each partition and the CS-DiD is re-estimated with the country as the panel unit (15 country clusters per partition); “HI ATT” and “LO ATT” are the resulting $\widehat{\text{ATT}}(k{=}{+}7)$ estimates and $\Delta = \widehat{\text{ATT}}^{\text{HI}} - \widehat{\text{ATT}}^{\text{LO}}$ is the Wald-type difference test (assuming independence of the two subsamples). “Stacked (30 c-p)”: the collapsed data are pooled across both partitions and the panel unit is set to the country–partition cell, yielding 30 unit-panel clusters and a single pooled ATT. Stars from two-sided Wald p-values: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.4 Robustness: additional elasticity specifications

Table 24 reports three additional elasticity specifications that are excluded from the main Table 4 for brevity. Columns (1) and (2) are no-cross specifications (excluding the business–technology interaction β_{AZ}) for the pre-2008 and post-2008 samples, respectively. Column (3) uses a first-differenced fixed-effects estimator (FD-FE) for the pre-2008 period, which avoids the equal-and-opposite artifact in the level-FE specification by directly estimating the effect of changes in intangible stocks. Across all pre-2008 specifications the business-intangible elasticity clusters between 0.36 and 0.40, and the post-2008 attenuation (to 0.12–0.20) is consistent across estimators.

Table 24: Additional elasticity robustness specifications

	(1) Level FE no cross 1995–2007	(2) Level FE no cross 2008–2017	(3) FD-FE no cross 1995–2007
Bus. intangibles (β_A)	0.390***	0.124**	0.398***
	(0.069)	(0.049)	(0.072)
Bus. intangibles$_{t-1}$	−0.398***		
	(0.072)		
Tech. intangibles (β_Z)	0.019*	−0.003	0.014
	(0.011)	(0.021)	(0.010)
Tech. intangibles$_{t-1}$	−0.014		
	(0.010)		
Tangibles (or Δ)	−0.411***	−0.247***	−0.398***
	(0.063)	(0.079)	(0.064)
Tangibles$_{t-1}$	0.395***		
	(0.067)		
Share foreign mkt. serv. (or Δ)	−0.772***	−0.995**	−0.902***
	(0.226)	(0.397)	(0.222)
Share foreign mkt. serv.$_{t-1}$	1.030***		
	(0.234)		
Observations	3173	3159	3173
Number of groups	312	318	312
Within R^2	0.202	0.170	0.199

Notes: Fixed-effects estimates with year effects and robust standard errors clustered at the value-chain panel identifier. Columns (1) and (2): log intangible stocks and their lags, no business–technology interaction. Column (3): first-differenced regressors (FD-FE); "Bus. intangibles" denotes $\Delta\ln(\text{bus. intangible stock})$; no lagged terms. The dependent variable is annual VC TFP growth in all columns. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.5 Model parameters: identification and calibration

The model's parameters are of two kinds. Seven are *estimated*: the minimum-distance procedure of Section 6.1 chooses θ_E^A, θ_E^Z, θ_C^A, θ_C^Z, μ_A, $\bar{T}_E$ and $\bar{T}_C$ to match the seven empirical moments, and Panel A of Table 5 reports the result. The remaining parameters are *calibrated*: they are fixed before estimation, either because they set the model's units, because they take conventional values from the quantitative macro and trade literatures, or because we impose them on external grounds. Panel B reports these. This appendix discusses identification for the first group and the basis of each value for the second.

Identification of the estimated parameters. The counts are equal, so the order condition is satisfied; we do not verify local identification formally, and treat the re-calibrations in Appendix A.7 as the practical check on the parameters that are least tightly pinned down.

Each moment moves one or two parameters most strongly. The pre-liberalization GDP ratio pins the relative size of the two regions, and so the ratio $\bar{T}_E/\bar{T}_C$. The two event-study TFP responses pin the business-capability elasticities and the corresponding size shifters: the European response identifies θ_E^A and $\bar{T}_E$, the upstream response θ_C^A and $\bar{T}_C$. The European business-capability response and the interaction elasticity β_{AZ} jointly identify μ_A: the first is a level change in A_E along the liberalization step, the second a slope interaction in the panel regression, and the cross-derivative of $\ln T_E$ with respect to $\ln A_E$ and $\ln Z_E$ is zero when $\mu_A = 0$. The business elasticity β_A maps to θ_E^A, which in the model equals β_A up to general-equilibrium feedback. Finally, the technology elasticity β_Z is what separates θ_E^Z from θ_C^Z: the model delivers a negative β_Z only when $\theta_C^Z > \theta_E^Z$, because European technology then raises China's TFP disproportionately through the cross-border bundle, while the levels of the two elasticities are further constrained by the two TFP responses. Because the marginal elasticities implied by (11) depend on the units of K_i^A and S_i, μ_A is identified only conditional on the model's normalizations.

Calibration of the predetermined parameters. The calibrated values fall into three kinds. Standard macro and trade parameters take conventional annual values: $\beta = 0.96$ and $\delta_K = 0.06$, which imply a required gross return of $\bar{R} \approx 0.102$; a capital share $\alpha = 0.35$, within the range the EU KLEMS accounts report for the market sector; and an Armington elasticity $\sigma_M = 4$.

Normalizations set the model's units. Europe is the numeraire for every level: $L_E = 1$, $\bar{a}_E = 1$ and $z_E = 1$, so China's endowment and shifters and both size shifters $\bar{T}_i$ are measured relative to Europe. The stock elasticity $\eta_A = 1$ is a maintained restriction rather than a pure normalization, since the marginal elasticity in (11) is $\eta_A + \mu_A \ln S_i$; setting it to one imposes a unit-elastic benchmark, so capability inherits the units of the stock.

The remaining values are imposed on external grounds and are the ones worth stating explicitly. China's labour endowment is $L_C = 4L_E$, roughly the ratio of China's manufacturing workforce to the EU15's over our sample; this is not a free normalization, because delivered scale enters the capability block. The depreciation rate of business intangibles is $\delta_A = 0.20$, at the lower end of the range used for organizational capital in the intangibles literature (Corrado et al., 2022); it matters because it converts the investment flow into the stock that enters (11), and re-calibrating at $\delta_A = 0.30$ leaves the objective and the sign of every result unchanged. Business-intangible investment is a fixed share of final output, $s_E^A = 0.06$ and $s_C^A = 0.03$: the European value lies within the range the EU KLEMS & INTANProd accounts report for business- and organization-related intangible investment (Bontadini et al., 2023; Corrado et al., 2022), while the Chinese value—for which no comparable accounts exist—imposes an intensity half Europe's. The substitution elasticities in the intangible bundles are set symmetrically at $\sigma_A = \sigma_Z = 6$, so that any asymmetry in the results comes from the estimated elasticities rather than from the bundle structure. The diffusion barriers are $\tau_A = 1.20$ and $\tau_Z = 1.10$; the ordering $\tau_A > \tau_Z$ reflects that codified technical knowledge travels across borders more readily than organizational practice, which tends to travel with the firm rather than the location (Keller, 2004; Bloom et al., 2012). At the calibrated equilibrium the domestic component accounts for 52 % of each region's technology bundle and 72 % of Europe's business bundle. Finally, $\bar{a}_C = 0.65$ places China's baseline organizational capacity below Europe's, and $z_C = 1$ makes baseline technology capacity symmetric. Appendix A.7 re-calibrates the model across ranges of z_C and of the diffusion barriers, and explains why z_C in particular has to be imposed.

A.6 Additional model moments and comparative statics

This appendix reports comparative statics that complement Section 6. The Step 0 to Step 1 changes in the non-targeted variables are already shown in Table 7.

Table 25: Additional Step 2 comparative statics: fragmentation around Step 1

	$\tau_{M,1}$+0.01	$\tau_{A,1}$+0.01	$\tau_{Z,1}$+0.01
EU TFP (T_E)	−0.236	−0.754	+0.046
CN TFP (T_C)	−0.069	−0.362	−0.208
EU output (Y_E)	−0.591	−1.113	+0.042
CN output (Y_C)	−1.253	−0.682	−0.239
EU business cap. (A_E)	−0.245	−0.782	+0.064
CN business cap. (A_C)	−0.196	−1.026	+0.073
EU scale (S_E)	−0.599	−1.148	+0.065
CN scale (S_C)	−0.994	−0.592	−0.297

Notes: Each column reports the percent change of endogenous variables relative to the Step 1 steady state when the indicated Step 1 exogenous object is perturbed by the stated amount and the economy is re-solved. Trade costs are iceberg. The results continue to show that higher goods-trade frictions and tighter diffusion barriers are contractionary for both regions. Higher technology-diffusion barriers (τ_Z) reduce China's TFP but slightly raise Europe's: because baseline technology capacity is symmetric, the barrier shrinks both effective bundles by the same proportion, so the difference in the TFP response reflects the conversion elasticities alone.

Table 25 then reports a small set of additional Step 2 comparative statics around the Step 1 equilibrium. The focus is on shocks that can be interpreted as partial geoeconomic fragmentation: higher trade costs in manufacturing goods (τ_M), in business-intangible services (τ_A), and in technology services (τ_Z). These exercises are best read as local comparative statics around the post-liberalization steady state. They are informative about the model's mechanisms, but not about a unique decomposition of the post-2008 slowdown.

The pattern in Table 25 is consistent with the discussion in the main text. Increases in manufacturing or business-services trade costs reduce output and TFP in both regions by shrinking the scale of cross-border production. The business-services margin (τ_A) remains especially relevant for the capability block, because higher τ_A lowers the effective business bundle directly. Higher technology-diffusion barriers (τ_Z) are the exception: a 0.01 increase in τ_Z reduces China's TFP by 0.21 % while *raising* EU TFP by 0.05 %. Because baseline technology capacity is symmetric, the barrier shrinks the two effective bundles by exactly the same proportion (0.43 % each), so the entire difference in the TFP response is attributable to the conversion elasticities θ_i^Z: China loses far more from the same reduction in access to foreign knowledge, and Europe gains on net through the resulting improvement in its relative competitive position. Tighter technology-diffusion barriers are thus the one fragmentation margin that is not contractionary for Europe in this model—the mirror image of the result that technology improvements benefit China disproportionately. These exercises are disciplined comparative statics rather than a literal account of any single post-2008 shock.

A.7 Robustness of the calibration

This appendix asks how much of Section 6.2 survives changes to the model's more discretionary elements. The technology-elasticity gap, and with it the result that a European technology improvement lowers European TFP, holds throughout; the size of the gap is more sensitive to these choices than its direction.

The cross-elasticity β_{AZ}. Spec B drops β_{AZ} from the objective and targets the business and technology elasticities from the no-cross specification, whose technology coefficient is *positive* (+0.019). Even so the calibration returns $\theta_C^Z \gg \theta_E^Z$, so the gap is not produced by targeting the cross-elasticity. Its nominal ratio is inflated by two corner solutions and should not be read as a magnitude.

The stock–scale interaction μ_A. Spec C fixes $\mu_A = 0$ and re-estimates everything else. The calibration then fails on the pre-liberalization GDP ratio (model 0.01 against a target of 3.0, objective 4.75 against 0.386). The reason is structural: with $\mu_A = 0$ and $\eta_A = 1$, (11) collapses to $a_i = \bar{a}_i K_i^A$, so capability is proportional to investment in both regions and the only remaining cross-region difference in the capability block is the exogenous shifter $\bar{a}_i$. The interaction term is the model's one endogenous source of cross-region variation in how investment maps into capability. It does not, however, generate the asymmetry: the capability feedback works against Europe in both technology experiments rather than amplifying the elasticity gap.

The business-capability elasticity θ_E^A. Because β_A is the one moment the model misses, one might worry that the optimizer inflates θ_E^A to close that gap and that the technology asymmetry is a by-product. Table 26 re-calibrates with θ_E^A held fixed at progressively lower values. Restricting it *raises* the ratio θ_C^Z/θ_E^Z, from 14.7 to 49.8, while the objective roughly doubles. The asymmetry at the estimated θ_E^A is thus a conservative value rather than an artifact, though at the most restrictive grid point θ_E^Z reaches its lower bound.

Table 26: Technology asymmetry robustness: sensitivity to θ_E^A

θ_E^A (fixed)	0.40	0.60	0.964 (baseline)
θ_E^Z	0.010	0.020	0.037
θ_C^Z	0.498	0.479	0.540
θ_C^Z/θ_E^Z	49.8	24.4	14.7
Objective	0.798	0.505	0.386
θ_E^Z at its lower bound	yes	no	no

Notes: Each column re-calibrates the model in MATLAB with θ_E^A held fixed and all remaining parameters estimated freely against the same seven moments; $z_C = 1$ throughout. The $\theta_E^A = 0.80$ cell is omitted because its solve failed. Restricting θ_E^A does not weaken the technology asymmetry — the ratio rises as θ_E^A falls — but the fit deteriorates sharply, the objective roughly doubling by $\theta_E^A = 0.40$, and at that point θ_E^Z reaches its lower bound.

Technology capacity z_C and the diffusion barriers. Two predetermined values deserve separate attention because the model gives no guidance on either. China's baseline technology capacity z_C is the special case: it is a level scale that enters *both* regions' technology bundles—China's directly and Europe's through the spillover term z_C/τ_Z—so it shifts the height of the bundles without changing how either region converts technology into productivity. It is also not identified: when estimated it runs to whatever upper bound is imposed. We therefore set $z_C = z_E = 1$, which takes no stand on which region holds the larger stock of technological knowledge and leaves the entire cross-region asymmetry to the estimated elasticities. A large z_C would instead build in a strong prior: at $z_C = 5$ only about a fifth of Europe's effective technology bundle would be domestic, which is hard to reconcile with a period in which EU15 research spending substantially exceeded

China's. The diffusion barriers τ_A and τ_Z are the model's only source of home bias, and their levels are likewise imposed.

Table 27 re-calibrates along both dimensions. Panel A varies z_C over a tenfold range: the elasticity ratio is essentially flat, between 14.1 and 15.6, and the European TFP response to a European technology improvement is negative at every value. What z_C absorbs is the split of the level between $\bar{T}_i$ and z_i, not the conversion elasticities—$\bar{T}_E/\bar{T}_C$ moves from 11 to 23. Panel B raises the barriers, which strengthens home bias at almost no cost in fit but weakens the spillover through which a European improvement reaches China; the model then needs an ever larger elasticity differential to reproduce the same β_Z, and at the two widest settings θ_E^Z reaches its lower bound. This is why the baseline keeps the barriers modest.

Table 27: Sensitivity to the predetermined technology capacity and diffusion barriers

Panel A. Predetermined technology capacity z_C (barriers at baseline)

	$z_C = 0.5$	Baseline ($z_C = 1$)	$z_C = 2.0$	$z_C = 5.0$
θ_E^Z	0.037	0.037	0.035	0.040
θ_C^Z	0.527	0.540	0.550	0.564
θ_C^Z/θ_E^Z	14.2	14.7	15.6	14.1
$\bar{T}_E$	1.234	1.228	1.213	1.172
$\bar{T}_C$	0.111	0.093	0.074	0.050
Objective	0.390	0.386	0.381	0.372
EU tech shock: ΔT_E	−0.053	−0.044	−0.034	−0.016
EU tech shock: ΔT_C	+0.294	+0.232	+0.168	+0.099

Panel B. Diffusion barriers (τ_A, τ_Z) (holding $z_C = 1$)

	Baseline (1.20, 1.10)	(1.50, 1.30)	(2.00, 1.70)	(2.50, 2.10)
Domestic share of Z_i	52%	55%	61%	65%
θ_E^Z	0.037	0.026	0.010	0.010
θ_C^Z	0.540	0.647	0.919	1.130
θ_C^Z/θ_E^Z	14.7	25.2	92.8	113.8
Objective	0.386	0.384	0.378	0.388
θ_E^Z at its lower bound	no	no	yes	yes

Notes: Every column is a full re-calibration of all estimated parameters against the same seven moments, run in MATLAB with three starting values per cell and the best retained. The EU tech shock rows report the percentage change in each region's TFP from a 1% increase in z_E at the Step-1 equilibrium.

Panel A: the technology-elasticity ratio is essentially flat across a tenfold range of z_C (14.1 to 15.6), and the European TFP response to a European technology improvement is negative at every value. What z_C absorbs is the level split between $\bar{T}_i$ and z_i.

Panel B: raising the barriers strengthens home bias at almost no cost in fit, but it also weakens the cross-border spillover through which a European technology improvement reaches China, so the model needs an ever larger elasticity differential to reproduce the same β_Z. At the two widest settings θ_E^Z sits on its lower bound of 0.01, so the ratio there is a corner artifact; the gap itself is present in every cell.

ETLA

Elinkeinoelämän tutkimuslaitos

ETLA Economic Research

ISSN-L 2323-2420
ISSN 2323-2420 (print)
ISSN 2323-2439 (pdf)

Publisher: Etlatieto Oy

Tel. +358-9-609 900
www.etla.fi
firstname.lastname@etla.fi

Arkadiankatu 23 B
FIN-00100 Helsinki