The Impact of Data Transfer Restrictions on Data Centre Geography

Otto Kässi[*1,2] and Vili Lehdonvirta[1]

[1]Aalto University
[2]Etla Economic Research

January 21, 2026

Abstract

National restrictions on cross-border personal data transfers potentially reshape where computation takes place, yet evidence on their impact on digital infrastructure remains limited. This paper studies how tightening data transfer regimes affects the geography of data centre investment. We construct a new country–year panel of global data centre openings and combine it with a novel classification of national data transfer laws. According to our findings, in the full sample, the estimated average effect of regulatory tightening on domestic data centre entry is positive but the estimates are imprecise. The increase is concentrated in larger economies and in countries with more downstream production structures, where the estimates are statistically significant. Event-study estimates show no systematic pre-trends, and regional placebo tests provide no evidence of cross-country spillovers within the region. A falsification test using cryptocurrency data centres yields null effects. To reconcile these findings with prior evidence that privacy regulation reduces overall compute demand, we develop a simple model in which regulation operates as a compliance quota that forces a minimum share of compute to be performed domestically. The model shows that tighter regulation can simultaneously reduce total compute use while increasing local compute and infrastructure investment, and it reproduces the observed

[*]We thank Linh Ngo for her excellent work in collecting and analysing the law dataset. This work is supported by an ERC grant (GEOCLOUD, 101141671).

heterogeneity patterns. Overall, the results suggest that tighter data-flow regulation shifts data centre entry toward domestic locations, consistent with a compliance-quota mechanism.

JEL F14, L51, L86, O33, D22

Keywords: Cross-border data flows; data localisation; data transfer restrictions; data centres; digital infrastructure; regulation; investment; compliance constraints; cloud computing; geography of computation; GDPR; weaponised interdependence; digital trade

1 Introduction

Many countries now regulate the cross-border flow of personal data. A large empirical literature studies how these rules shape *intangible* outcomes such as innovation, productivity, digital trade, and firm behaviour in data-intensive activities. Much less is known about how data-flow regulation maps into *physical* investment. This gap matters because digital production still relies on fixed, place-based capital. Data centres are the factories of the digital age - they bundle servers, cooling, and specialised building infrastructure to deliver storage and computing services, and their location determines where computation takes place.

This paper links data-flow regulation to the geography of data centre investment. The key mechanism is a composition effect: when cross-border transfers become restricted or costly, firms can comply by relocating the processing of affected data to domestic infrastructure, increasing the domestic share of compute even if the underlying business activity is unchanged. But regulation can also change the level of activity. GDPR-style privacy rules have been shown to reduce output and raise costs in data-intensive markets, which can lower total demand for compute (Janßen et al., 2022; Demirer et al., 2024; Chang et al., 2023). These two forces work in opposite directions for domestic buildout: localisation raises the domestic share of compute, while reduced activity lowers the overall scale of compute. The net effect on domestic data centre entry is therefore theoretically ambiguous, and is ultimately an empirical question.

We ask: do tighter cross-border data transfer regimes increase domestic data centre entry? To answer this, we combine two new pieces of measurement. First, we construct a country-year panel of global data centre openings from S&P Global Capital IQ Pro Real Estate, which tracks property-level data centre facilities across 101 countries using a globally standardised methodology. Second, using LLM-based classification, we build a harmonised panel measure of national cross-border data transfer regimes by classifying each country's legal framework into a restrictiveness scale using the full legal text, with targeted manual validation. In the empirical analysis, we focus on discrete regime changes, in particular the first adoption of laws at least as restrictive as the GDPR.

As of 2024, 172 countries and autonomous regions had enacted laws or rules that restrict or complicate the transfer of personal data abroad.[1] These policies are often justified in terms of privacy, sovereignty, or security, yet they may also alter

[1] These rules are catalogued in a regularly updated format in Greenleaf (2021). The latest update at the time of writing is Greenleaf (2025). See also Ferracane (2017).

the geography of computation. Despite the centrality of this issue for digital trade and production, systematic evidence on how data transfer restrictions influence local data centre investments across countries remains limited.

We implement a two-way fixed effects event-study design around regulatory tightenings. We find that tightening cross-border data transfer restrictions is followed by higher domestic data centre entry on average, although the full-sample estimate is imprecise and not statistically significant. The effects are heterogeneous: responses are larger in countries with higher GDP and population, and in countries whose industries are closer to final demand (i.e., more *downstream*; Antràs et al., 2012).

This finding runs counter to a body of research concentrating on the GDPR - in particular Demirer et al. (2024) and Chang et al. (2023) - which find that the GDPR resulted in lower compute demand for the affected firms. To reconcile these results, seemingly contradictory findings, we develop a simple theory model with a compliance quota. Under this model, tighter regulation can both decrease total compute demand but increase the demand for domestic compute and build-out of local data centre capacity. We further show, that this model reproduces the comparative statics we observed empirically regarding GDP, population and downstreamness.

Finally, our model implies, that the investment response depends on the relative costs of local and remote compute. In the model, tightening results in more data centre entry, when the price differences between remote and local compute are such that forced reallocation to local compute does not result in large unit price increase of compute. Consistent with this, we find that the post-tightening increase in domestic data centre buildout is larger in countries with cheaper local electricity.

Our paper connects to several literatures. First, it relates to work on how trade policy affects domestic investment (e.g., Pierce and Schott, 2018; Amiti et al., 2020; Handley and Limao, 2015; Handley and Limão, 2017; Caldara et al., 2020). A common finding in this literature is that raising trade barriers can reduce investment by increasing input costs and policy uncertainty, particularly for firms exposed through global value chains. At the same time, when policy restricts access to a close foreign substitute, it can induce substitution toward domestic choices in the affected goods. We document this substitution in the context of an intangible good: tighter cross-border data transfer rules restrict access to foreign compute for certain data, and are followed by higher domestic investment in data centre capacity.

Second, we contribute to a growing empirical literature on the economic effects of the GDPR and related privacy regulation (e.g., Peukert et al. (2022); Frey and Presidente (2024); Demirer et al. (2024); Goldberg et al. (2024); Aridor et al. (2020); Johnson et al. (2023); Janßen et al. (2022); Koski and Valmari (2020); Gupta et al. (2022); Sisto and Van der Marel (2025); Zhang et al. (2025); Jia et al. (2025); Sun and Trefler (2023)). We view the GDPR as one prominent instance within a broader set of cross-border data transfer regimes and study domestic investment responses across countries using a globally comparable panel.

Third, we contribute to a set of papers that study data as a intangible factor of production, including Farboodi and Veldkamp, 2021; Corrado et al., 2022; Agrawal et al., 2019; Goldfarb and Tucker, 2019. Our model formalises the complementarity between intangible data and physical compute and adds a location choice between local and remote compute that is distorted by regulation.

Fourth, our paper links to small but growing literature concerning the 'geography of cloud computing', which studies the determinants of data centre locations (Pan Fang and Greenstein, 2025; Lehdonvirta et al., 2025; Bonfiglioli et al., 2025; Tian et al., 2025).

Finally, we relate to the "weaponised interdependence" argument in Farrell and Newman (2019), which discusses how some states exploit their jurisdiction over key parts of the global digital infrastructure such as data centres to monitor or coerce foreign actors. Our contribution is to show, using cross-country evidence, that data transfer restrictions are followed by a reallocation of data centre investment toward domestic locations, consistent with the idea that laws governing data flows can be used to limit exposure to foreign jurisdiction over where data are processed.

2 Data

2.1 Laws regulating the cross-border flow of data

Laws that regulate cross-border data flows typically specify the legal and procedural requirements for sending data abroad, such as adequacy decisions, contractual safeguards, consent-based derogations, or government approvals (Casalini and López González, 2019). For the purposes of this study, we focus on these cross-border transfer restrictions for *personal data*, and treat them as a measure of the practical difficulty of processing personal data abroad.

We interpret stricter regimes as expanding the set of data and use-cases for which cross-border is illegal, uncertain, or operationally costly enough that firms choose (or are forced by law) to process data within their borders. Accordingly, ranking countries' laws by overall restrictiveness provides a consistent measure of how binding the domestic-processing requirement is in practice, and thus of the implied increase in the minimum domestic share of compute associated with regulated data.

Our dataset on national data governance frameworks extends the cross-border data regulation dataset compiled by Greenleaf (2021). It documents the primary national data privacy laws that govern cross-border data transfers and the domestic storage of personal data. For each country, we identify the main comprehensive privacy legislation concerning personal data, recording the introduction of new frameworks or replacements of existing ones, and their dates of enactment. Minor amendments that do not alter the overall structure of the legal framework are excluded. This allows us to construct a balanced country–year panel of data governance regimes. We focus on economy-wide ("comprehensive") personal-data frameworks rather than sectoral rules (e.g., health, finance) because the former are the relevant margin for broad data-centre investment decisions and for cross-industry comparability.

While Greenleaf (2021) provides a comprehensive inventory of relevant laws, it does not assess their restrictiveness. We extend the dataset by classifying each national framework using the taxonomy of cross-border data flow regulation developed by Casalini and López González (2019). This taxonomy distinguishes seven levels (0–6) of regulatory restrictiveness, ranging from the absence of specific cross-border provisions (0) to frameworks requiring case-by-case authorisation for each transfer (6). Intermediate levels correspond to conditional regimes—such as accountability-based systems (1), firm-assessed adequacy regimes (2), and public-authority-based adequacy regimes with contractual or procedural safeguards (3–4). We note, that this taxonomy ranks regimes by a single dimension, the practical difficulty of transferring personal data abroad. Thus, it makes no claims about the comparability of different regal regimes in their objectives or institutional design.

To assign each country–year observation to one of these levels, we employ a large language model classifier that analyses the full legal text. The model is prompted with structured criteria derived from the taxonomy definitions in Casalini and López González (2019). The resulting classification provides a panel measure of data flow restrictiveness that is conceptually aligned with the OECD framework and enables consistent cross-country and temporal comparisons. We present the classification of laws with examples in Table 1.

Classifying legal texts inevitably involves a degree of subjectivity. We address this through two validation steps. First, we manually review a large sample of laws to confirm consistency between human and LLM-based classifications. Second, in our empirical analysis, we use a binary indicator that equals one in the first year a country enacts a regime at least as strict as the GDPR. This dichotomisation reduces sensitivity to small ranking errors between intermediate levels. Remaining (classical) measurement error would bias estimated effects toward zero, making our results conservative.

The staggered adoption of data transfer laws is visualised in 1. Our “treated” units consist of 63 countries that implemented a law restricting cross-border data transfers.

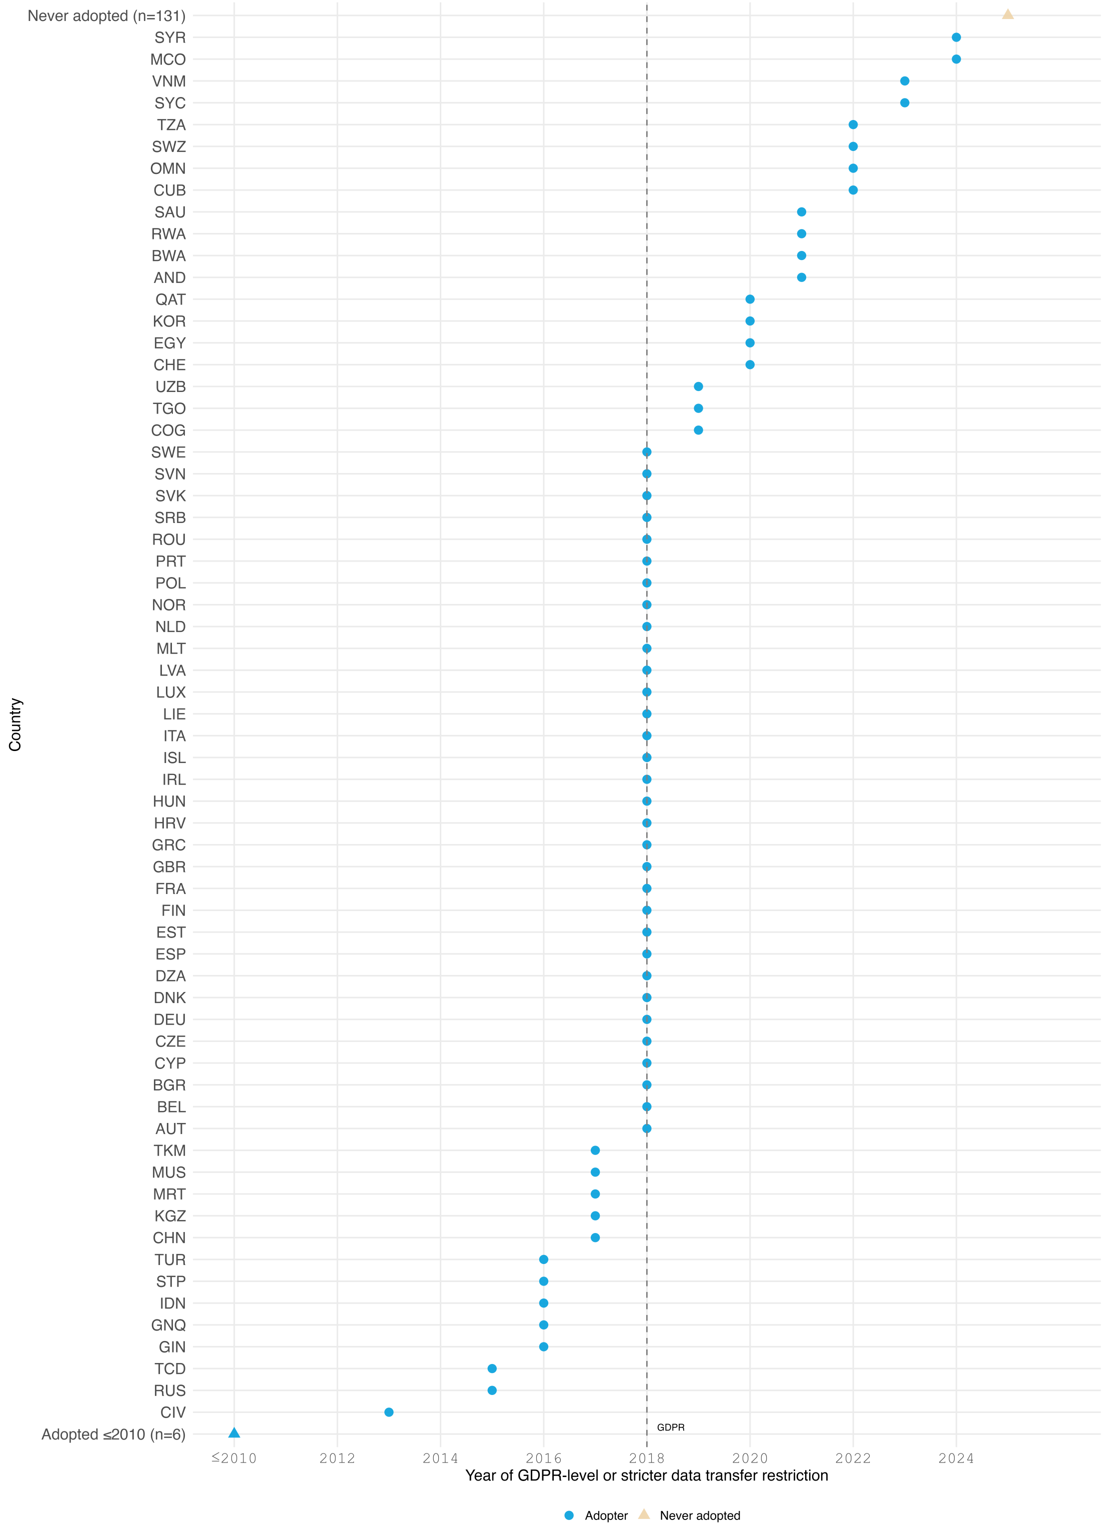

Figure 1. Staggered adoption of laws regulating data transfer between years 2010-2024

OECD High-level Category	Sub-Level (0–6)	Description	Example Countries
No Regulation	0. No Regulation	No domestic rules governing cross-border flows.	Afghanistan, Bangladesh
Ex-Post Accountability	1. Free-Flow with Ex-Post Accountability	Transfers free; exporter liable if data misused abroad.	Canada, Hong Kong
Conditional on Safeguards	2. Self-Assessment of Adequacy + Fallbacks	Exporter judges adequacy; otherwise relies on consent, contracts, or necessity.	Australia, India, Mexico
	3. Public-Assessment of Adequacy + Options	Public authority certifies "adequate" destinations; others need SCCs/BCRs/consent.	Japan, United States
	4. Public-Assessment + Options + Processing Conditions	Same as 3 plus requirement that foreign processing apply sender's rules; exporter liable downstream.	European Union, United Kingdom
Conditional on Ad-Hoc Authorization	5. Case-by-Case Authorization with Adequacy Fallback	Transfers to "adequate" jurisdictions flow freely; others require regulator approval.	Russia, Saudi Arabia
	6. Case-by-Case Universal Authorization	All transfers require prior regulator approval; often tied to localization.	China, Egypt

Table 1. OECD Cross-Border Data Transfer Restrictiveness

2.2 Data centres

The dependent variable measures new data centres built in a given country–year. These come from the S&P Global Capital IQ Pro Real Estate, which is a commercial database of real estate related data. According to S&P, their dataset covers 95,000 properties across 101 countries. Capital IQ analysts draw on public company disclosure data such as annual and quarterly financial reports, regulatory filings, investor presentations and other market documents. The main strength of the dataset for our analysis is that it is standardised to be globally consistent. Its methodology allows us to map data centre construction consistently across jurisdictions and time.[2]

The coverage of the dataset is extensive but not exhaustive. It is strongest for listed firms and large private operators that disclose property-level information through regulatory filings or investor communications. Coverage is therefore biased toward jurisdictions with robust financial disclosure requirements and toward larger commercial or institutional properties, including hyperscale data centres owned by public or major private firms. Facilities owned by smaller domestic operators or entities not subject to public reporting, such as municipal or sovereign facilities, small colocation providers, or data centres embedded in enterprise campuses, may be under-represented or missing.

The pattern implies, that our data is oriented towards more formal, capital-intensive segment of the market, while omitting smaller scale or informal capacity. If the investments respond similarly, or more strongly, to restrictions on cross-border data flows, our estimates of the policy effects would likely understate the total impact on domestic data infrastructure investments. Consequently, our results can be interpreted as conservative estimates of the overall impact of cross-border data transfer restrictions.

2.3 Additional country level data

In addition to the cross-border data transfer law and data centre variables described above, I use several country level characteristics to explore the heterogeneity in the effects of cross-border data flow restrictions. These variables capture structural differences in market size, development and industry composition that may have an impact on the investment response to the regulatory shock.

Population and GDP are obtained from the World Bank's World Development Indicators (WDI) database, using the series SP.POP.TOTL (total population), and NY.GDP.MKTP.KD (real GDP).

To proxy the differences in personal data intensity between countries, I use the downstreamness indicator by Mancini et al. (2024). They build on the theoretical framework from Antràs et al. (2012), where downstreamness measures how close to the final demand an industry is. Intuitively, the measure captures an average of how many production stages lay between the end user demand and an industry's production. A higher value indicates that an industry's production is closer to end users, while lower values correspond to upstream, intermediate-goods produc-

[2] This data has been extensively used in real estate economics and finance research (see, e.g., Ling, Wang and Zhou (2021); Ling, Naranjo and Scheick (2021)).

tion. We assume that industries closer to the end user produce and analyse more personal data compared to upstream industries.

In addition, we assemble two country-level infrastructure measures that will be used later to test cost- and connectivity-related predictions of the conceptual framework. First, we use total electricity generation (TWh) from the U.S. Energy Information Administration's international electricity statistics.[3] Second, we use international bandwidth from the International Telecommunication Union (ITU).[4] Both measures are taken in 2010.

We incorporate these measures as dichotomous variables, high/low, corresponding to above and below global median, measured at the start of our sample in 2010. Using measures from 2010 allows us to treat these characteristics as pre-determined with respect to subsequent changes in either cross-border data transfer regulation and local data centre investment. This reduces concerns related to reverse causality where regulation (or the induced data centre buildout) could mechanically raise GDP, shift production structure or affect population, and thereby contaminate the heterogeneity split. Due to relatively small sample sizes, we have opted to use a binary high/low split rather than a continuous measure.

2.4 Descriptive Statistics

This section reviews the descriptive statistics related to the data. We begin by describing the dependent variable, new sites by each year-country cell. Table 2 shows that most country-year cells have zero data centre openings, and their distribution is right-skewed. In the full sample, 73% of country–years have zero openings, while the top 1% of observations reaches 34 openings and the maximum is 189. Treated country–years have fewer zeros (58% vs. 76%) and substantially higher mean entry (3.44 vs. 1.49) and a heavier upper tail. We plot the cumulative sum of data centre sites for treated and non-treated countries in Figure 2.

We report the baseline means of the two groups in Table 3. Ever-treated countries are, on average, larger and more connected: they have higher 2010 GDP and population, and substantially higher international bandwidth, consistent with regulation being adopted earlier in bigger and more globally integrated economies. In contrast, downstreamness is very similar across the groups.

Table 4 presents a correlation matrix calculated from the binary baseline split indicators (defined using 2010 values). It shows that correlations across the splits

[3] https://www.eia.gov/international/data/world/electricity/electricity-generation

[4] https://datahub.itu.int/

are mostly low-to-moderate, indicating that the heterogeneity groups do not measure the same underlying dimension of the data. The main exception is that GDP and the infrastructure proxies are more tightly related, with richer economies having both greater energy supply and higher connectivity.

Treated	N	Share zero	Mean	Std. dev.	Mean (x>0)	P50	P75	P90	P95	P99	Max
Never treated (Treated=0)	2443	0.76	1.49	7.22	6.30	0	0	3	6	29	141
Treated country-years (Treated=1)	482	0.58	3.44	15.82	8.21	0	2	7	10	116	189
All country-years	2925	0.73	1.81	9.23	6.80	0	1	3	7	34	189

Table 2. Distribution of new data-centre site openings by treatment status.

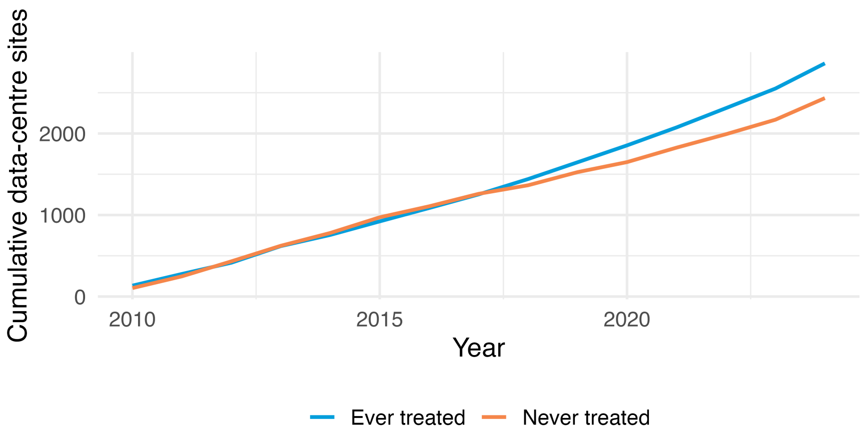

Figure 2. Cumulative number of data centre sites in ever-treated and never-treated countries.

Variable	N (Never)	Mean (SD) Never	N (Ever)	Mean (SD) Ever	Diff. means
Downstreamness (2010)	104	2.08 (0.36)	68	2.12 (0.38)	0.04
Electricity generation (2010, TWh)	119	87.83 (404.89)	66	157.09 (532.69)	69.26
GDP (2010 USD, bn)	123	278.18 (1532.63)	68	441.76 (1093.96)	163.57
International bandwidth (2010, Gbit/s)	120	198.80 (933.70)	67	467.93 (1064.49)	269.14
Population (2010, mn)	126	31.69 (117.68)	69	42.53 (162.94)	10.84

Table 3. Baseline (2010) country characteristics by treated status (ever-treated vs. never-treated). Cells report means, with standard deviations in parentheses.

3 Estimation method and the empirical model

We estimate the standard two-way fixed effects model

$$\Delta Y_{c,t} = \alpha_c + \eta_t + \beta D_{c,t}^{Tighten} + \varepsilon_{c,t}, \tag{1}$$

	GDP below median	Pop. below median	Downst. below median	Electricity gen. below median	Cross-border bandwidth below median
GDP below median	1	0.5	-0.11	0.76	0.69
Pop. below median		1.0	-0.20	0.41	0.26
Downst. below median			1.00	-0.12	0.01
Electricity gen. below median				1.00	0.67
Cross-border bandwidth below median					1.00

Table 4. Correlation matrix between the various heterogeneity splits.

where countries are indexed by c and years by t. $\Delta Y_{c,t}$ is the change in number of data centres in country c between years t and $t-1$. The dummy variable $D_{ct}^{Tighten}$ gets a value 1 the first time a country's data localisation law exceeds the level 4, or is at the level of GDPR or tighter.[5]

The global, staggered introduction of data localisation laws provides a useful source of quasi-experimental variation for studying their effects on local data centre investment. From the point of view of individual firms, the timing of a country's localisation law can be treated as plausibly exogenous, since firms do not determine regulatory timing. At the same time, country-level adoption may still correlate with broader trends in digitalisation or economic development. A two-way fixed effects event-study framework helps account for time-invariant country differences and common global shocks, allowing us to isolate the within-country changes associated with localisation laws.

This model comes with its own set of assumptions needed for causal identification. The key requirement is that, absent the adoption of localisation laws, treated and untreated countries would have followed parallel trends in data centre investments, conditional on time and country fixed effects. In other words, the timing of law adoption should be unrelated to unobserved country-specific shocks that also affect investment. In addition, the Stable Unit Value Assumption (SUTVA) should hold. This requires, that a localisation law in one country does not itself deter or induce data centre investments in other countries.

While both of these assumptions are ultimately untestable in an observational setting, we report pre-trend and placebo estimates that support them.

Another key challenge with staggered adoption designs is that the standard two-way fixed effects models compare units with different treatment timings, sometimes treating already-treated units as controls for those treated later. To address this, we follow Sun and Abraham (2021), who propose an event-study estimator that reweights group-time average treatment effects to avoid such "forbidden" comparisons. In this framework, countries that have not yet adopted a localisation law serve as valid controls for those that have, while already-treated countries are

[5] Choosing level 4 as the definition of treatment is motivated both by theoretical and practical reasons. First, it corresponds to a material tightening of transfer conditions. Although mechanisms such as Standard Contractual Clauses or Binding Corporate Rules remain legally available, their use typically involves added compliance cost, documentation, and regulatory scrutiny (European Parliament and Council of the European Union, 2016). In practice, Level 4 marks the point where transfers become conditional rather than routine. Additionally, from a more practical perspective, a substantial share of the treated countries - including the EU - in our data are at level 4, so setting the threshold higher would sharply reduce number of treated cohorts and statistical power.

excluded from the control group once they enter treatment.

Intuitively, the estimator decomposes the event-study coefficients into group-time average treatment effects $ATT_{g,t}$, defined for each cohorts of law adopters (g) and calendar year (t). Each $ATT_{g,t}$ is estimated using the corresponding not-yet-treated and never-treated countries as the control group. More formally, let G denote the country's first year of treatment ("cohort"). The outcome variable is the first difference of Y_{ct}, i.e. the number of new data centres opened in country c in year t. For each cohort g and calendar year t, we define the cohort-time treatment effect as

$$ATT_{g,t} = \mathrm{E}[\, \Delta Y_{gt}(1) - \Delta Y_{gt}(0) \,|\, G = g] \,, \tag{2}$$

where the expectation is taken over countries whose first treatment year is g and ΔY_{gt} denote potential outcomes for the first difference. $k = t - g$ is the event time. For each event time,

$$\beta_k = \sum_{g \in \mathcal{G}(k)} w_{g,k} \; ATT_{g,\, g+k}, \tag{3}$$

where $G(k)$ is the set of cohorts with observed event time k, and the weights $w_{g,k}$ are non-negative, sum to one, and depend on cohort size and the composition of not-yet-treated controls each year. Already-treated countries never contribute to the control group, so each $ATT_{g,t}$ is identified exclusively from never-treated and not-yet-treated units.

To back out the tretment coefficient β, we then calculate a cohort size weighted mean of the β_k's. This approach eliminates the contaminated comparisons arising when earlier-treated units act as controls for later treated ones.

4 Results

4.1 Main results

We start with the baseline TWFE estimates from 1, summarised in Table 5. In the full sample, the tightening of the data law is followed by roughly half an additional data centre per year relative to the pre-period. Relative to the pre-treatment mean, this implies an increase of about 35%.

Turning to analyses by subgroup, in the second and third column of Table 5. Industries positioned closer to end-users exhibit a stronger investment response (approx. 60% increase over the pre-treatment mean). The pattern is consistent with more customer-facing sectors being more exposed to regulatory frictions in cross-border data flows.

In columns 4 and 5, we repeat the analysis with two using data split into two groups by population size. We find, that the effect is statistically indistinguishable from zero in the low-population group of countries, while the effect is positive and statistically significant in the high-population group. or the latter group, the magnitude, about 1.2 additional data centres per year, amounts to a 31% increase over the group-specific pre-treatment mean.

Finally, in columns 6 and 7, we report the estimates from the model, where we have split the data by two by country GDP. Once again, the response is concentrated among the larger economies. Countries above the median GDP show an estimated effect of roughly 1.2 additional data centres annually, or around 40% relative to their pre-treatment mean. For poorer countries, the tightening appears to have no measurable impact.

	(1) Full sample	(2) Upstream	(3) Downstream	(4) Low population	(5) High population	(6) Low GDP	(7) High GDP
Estimate	0.834	0.583	0.608*	0.104	1.286.	-0.164	1.195***
Std. error	(1.247)	(2.060)	(0.239)	(0.112)	(0.698)	(0.170)	(0.339)
N countries	195	86	86	98	97	96	95
N observations	2925	1290	1290	1470	1455	1440	1425
Treated pre mean	2.308	3.449	1.005	0.544	3.824	0.459	3.020

Table 5. This table reports the effect of tightening data-flow restrictions on new data centre buildup. Estimates are obtained from a two-way fixed-effects model implemented with Sun–Abraham weights. Treatment time is defined as the first year in which a country imposes a cross-border data-flow restriction at least as stringent as the GDPR. Standard errors are clustered ont he country level. Significance levels: . = 10%, * = 5%, ** = 1%, *** = 0.1%.

4.2 Event study analysis

One potential concern for the TWFE analysis is that it might be driven by pre-tightening trends. These could either be driven by policy anticipation: if governments announce data transfer restrictions well advance before they come into force, operators may adjust their investment in anticipation of the policy, resulting in increase in data centre investments prior to the treatment. Moreover, the policy timing itself could be endogenous: countries experiencing a boom data centre boom for reasons unrelated to policy might be more likely to tighten cross-border data rules, which would result in pre-treatment growth and policy choices being correlated with one another. Both of these factors would threaten the TWFE identification strategy, which relies on measuring a weighted difference between not-yet treated and treated units' time trajectories.

However, we present evidence against this in the event study plots (Figure 3). Across all 7 TWFE specifications, we find that the pre-treatment are mostly around zero, with only a few pre-treatment estimates differing from zero at 5% risk level.

The event study plots also point to the relative uncertainty of the estimates. Even for the splits that have a positive average effect over the entire post-period, the individual year-effects are mostly zero.

4.3 Evidence against policy contamination across countries

Another threat to identification of the TWFE model are spillovers. If a policy in country c results in less investment in its other countries, this would result in a spillover from country c to other countries, which, in turn would result in violation of the Stable Unit Treatment Value Assumption (SUTVA), and consequently invalidate the research design.

A simple way to assess whether such spillovers are present is to ask, whether the neighbouring countries' outcomes systematically move when country c tightens its data transfer laws. The logic is the following: if the SUTVA holds, the potential outcomes of country $j \neq c$ must be unaffected by the treatment status of country c. This test is intentionally local: it is designed to detect displacement or complementarities to geographically proximate countries. It does not test for reallocation to distant locations outside the 'nearby' set (such as reallocation between the EU and the US).

Unlike the standard approaches that that capture spillovers by aggregating outcomes across neighbouring units (e.g. Di Tella and Schargrodsky (2004); Jardim et al. (2024)), we isolate average spillovers by constructing a leave-one-out outcome. Specifically, we define:

$$\Delta Y_{c,t}^{near} = \sum_{j \in r(c), j \neq c} , \tag{4}$$

where $r(c)$ denotes the region containing country c. If law tightening in c diverts investments away from nearby countries (or triggers complementary investments in them), the sum of neighbours outcomes will exhibit a systematic pattern around country c's treatment date. If SUTVA holds, $\Delta Y_{c,t}^{near}$ should remain approximately flat apart from noise.

We then estimate the Sun-Abraham TWFE event study model outlined in Equation (2) using $\Delta Y_{c,t}^{near}$ as the dependent variable. This produces an average

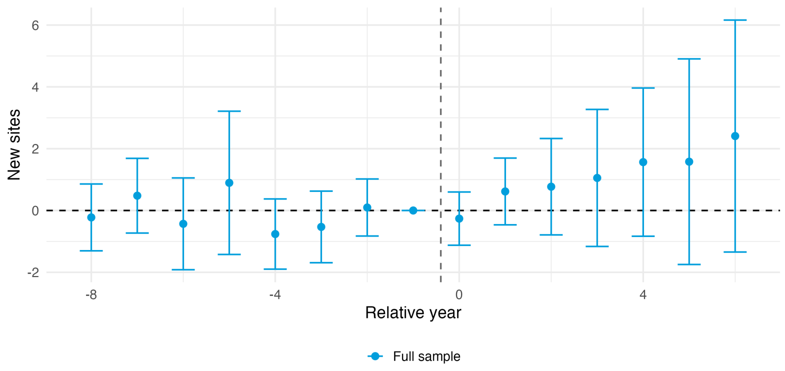

(a) Event study: full sample

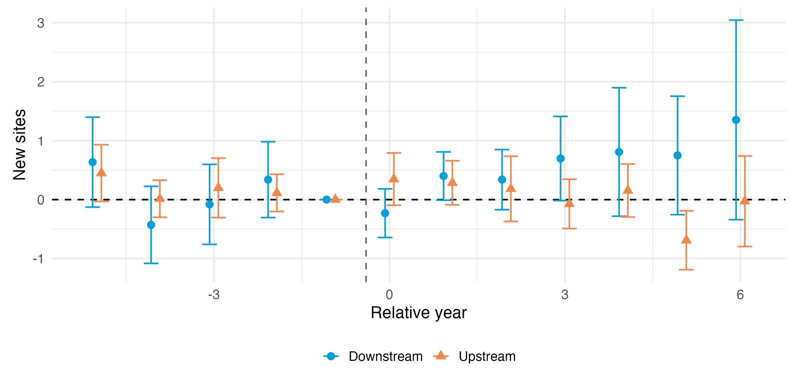

(b) Event study: Downstreamness split

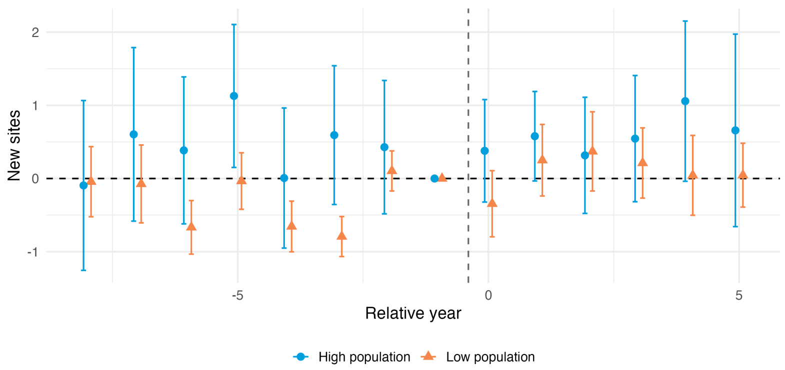

(c) Event study: population split

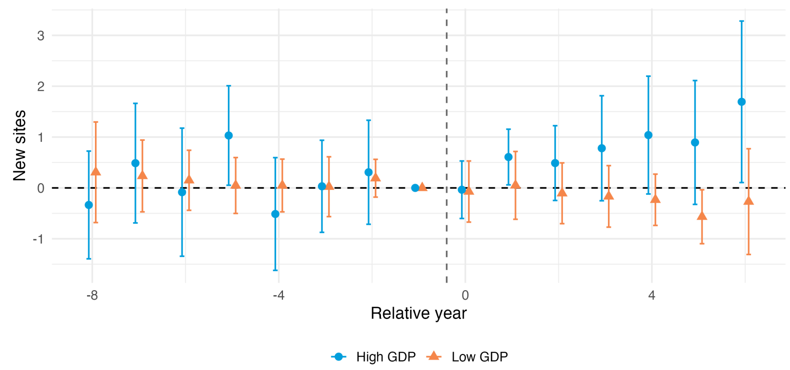

(d) Event study: GDP split

Figure 3. Event study pictures. The vertical lines correspond to 95% confidence intervals.

contamination profile: if the coefficients are approximately zero pre- and post-treatment, this is evidence against contamination.

We define "nearby" countries as countries that are in the same region $r(c)$ according to the 23 region classification used by World Bank, and others.[6]

By construction, a flat $\Delta Y_{c,t}^{near}$ rules out systematic within-region spillovers, but it remains silent with reallocation toward countries outside $r(c)$.

We do two versions of the spillover validation test. One with full sample and another with data winsorised at the 95% level. Results are shown in Figure 4.

The estimated coefficients are flat and centered around zero both before and after the law change, providing no evidence of systematic spillovers within regions as defined by $r(c)$. This supports SUTVA against local contamination, though it cannot exclude displacement to more distant hubs outside the region classification.

[6] We use the `countrycode` R package (Arel-Bundock et al., 2018). The regions are: South America, Caribbean, Central America, Northern America, Southern Asia, Western Asia, South-Eastern Asia, Eastern Asia, Central Asia, Middle Africa, Western Africa, Southern Africa, Eastern Africa, Northern Africa, Southern Europe, Western Europe, Eastern Europe, Northern Europe, Polynesia, Australia and New Zealand, Melanesia, and Micronesia.

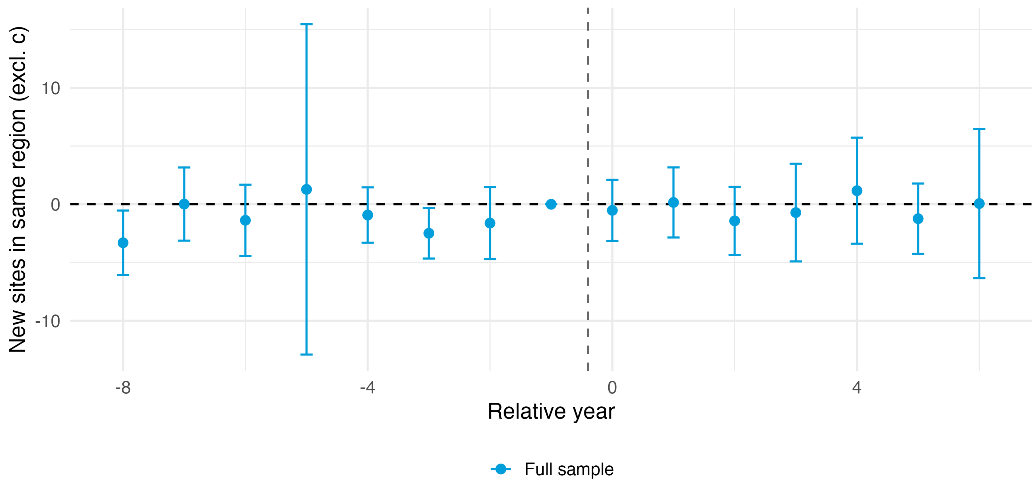

(a) Test for regional treatment spillover.

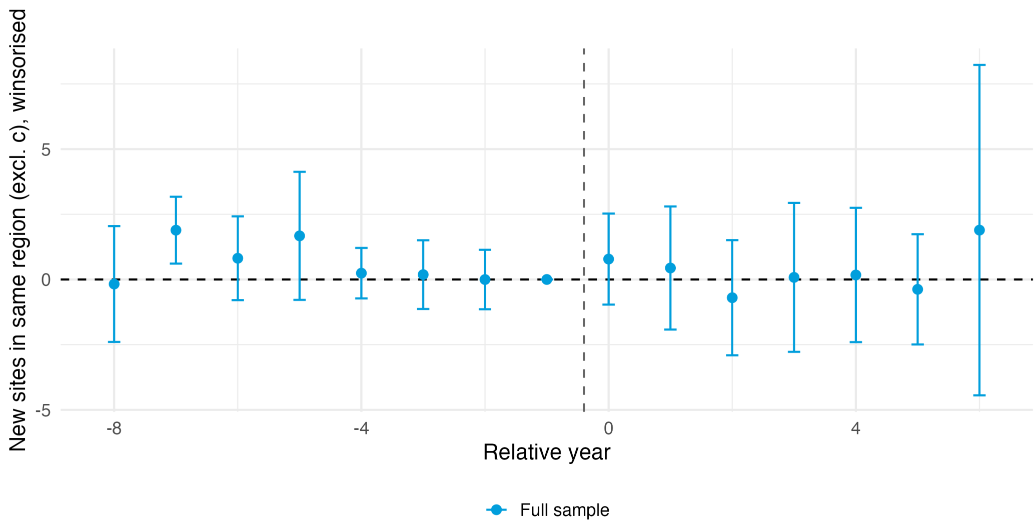

(b) Test for regional treatment spillover (winsorised sample).

Figure 4. Regional spillover tests.

4.4 Robustness check: only crypto data centres

As an additional falsification check, we show, that personal data transfer restrictions have zero impact on crypto data centre buildout. Cryptocurrency mining is structurally different from most other data processing as it does not rely on personal data. Consequently, they do not face compliance burdens under the the legal regimes we study.

We report the results of this exercise in Figure 5. Cryptocurrency data centres' location decisions primarily depend on energy and cooling costs, and infrastructure constraints, instead of data tranfer legislation. Finding no effect therefore increases confidence that our main estimates are not mechanically picking up general data centre buildup trends.

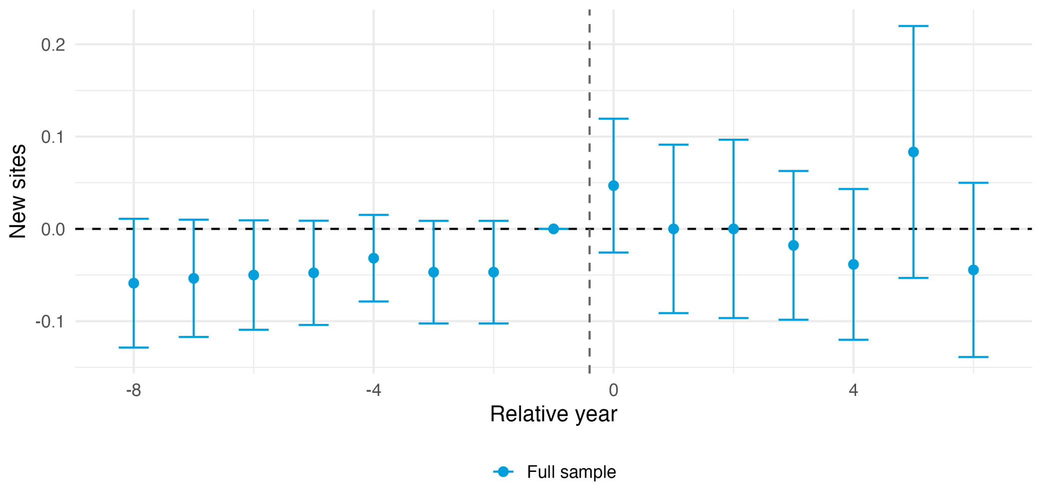

Figure 5. Impact of data transfer restrictions on crypto data centre sites.

4.5 Sensitivity of the results to outliers

One potential concern is that the results are based on conditional means of right-skewed count outcomes. Given the distribution of the outcome variables, it is at least theoretically possible that the estimation results are driven by a small number of extreme observations. To demonstrate that this is not the case, we rerun our estimates with the outcome variable winsorised at the 95% level.

The results are presented in Fig. 6. A comparison between the baseline estimates (in blue) and the winsorised estimates (in orange) reveals that the two sets of estimates are both quantitatively and qualitatively similar. This indicates that the main results are not driven by outliers. The only exception is the full-sample estimate, where winsorising reduces the standard error while leaving the point estimate largely unchanged, suggesting that a small number of extreme observations mainly affect precision rather than identification.

4.6 Sensitivity of the results to estimator choice

To demonstrate that the results are not driven by the choice of estimator, we re-estimate the model using the estimator proposed by Callaway and Sant'Anna

(2021). This estimator provides an alternative approach to Sun and Abraham (2021) for estimating treatment effects in TWFE settings with staggered treatment adoption.

The resulting estimates are also reported in Fig. 6. Across all subsamples, the Callaway–Sant'Anna estimates (in purple) are close to the corresponding Sun–Abraham estimates in terms of point estimates, but have systematically larger confidence intervals in several cases.

This pattern is consistent with the small-sample properties of the two estimators. In our setting, sample sizes and treated cohorts are mostly small, which reduces the the effective sample size available. The Sun–Abraham estimator often achieves higher efficiency, while delivering estimates that are quite similar to those obtained from the more flexible Callaway–Sant'Anna approach, but with smaller standard errors. This builds confidence that our main results are not driven the estimator choice.

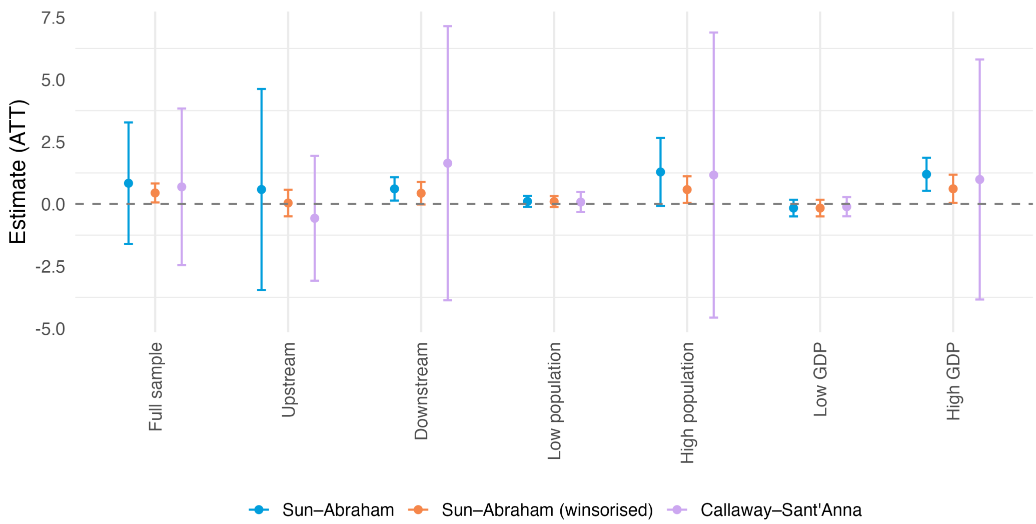

Figure 6. Caption describing the image.

5 Conceptual Framework

After documenting that tighter legal restrictions on cross-border data transfers are followed by higher local data centre investment, we introduce a simple model that is able to generate the comparative statics that we estimated from the data.

The model combines a standard production technology with compute and data as complements and a compliance requirement that forces a minimum domestic share of compute. This setup can reproduce the main empirical patterns from the previous section.

A common approach in macroeconomics and international trade is to represent policy and regulatory frictions as "wedges" that distort firms' choices (Chari et al., 2007; Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009). Particularly, in the data-economy context, recent papers such as Demirer et al. (2024) and Chang et al. (2023) model regulation as a wedge that lowers the marginal productivity of computational resources and data. A general implication of this assumption is that such wedges reduce the use of the regulated input and, through complementarity, also reduce the use of its complements. Our estimates point in the opposite direction: in some countries, tighter regulation is followed by more data centre investment.

We take an alternative route. Following (Kwon and Chun, 2009), we model regulation as a compliance quota. Rather than imposing an explicit cost on cross-border data transfers, we require that a minimum share of compute be performed domestically. Because regulated data must remain within the country, compliance can be interpreted as forcing the associated compute to take place within the country. Under this interpretation, the policy index τ is a reduced-form measure of regulatory scope. Higher corresponds to a broader set of data for which cross-border processing is infeasible, uncertain, or administratively burdensome enough that firms treat it as effectively non-transferable. $\omega(\tau)$ is the share of total compute associated with regulated data. Next section demonstrates that this model can theoretically generate the empirical patterns from the previous section.

5.1 Model

A representative firm in country c uses a strictly positive amount of compute C and data D to produce a digital good for the local market:

$$Y_c = C^{\alpha} \log(1 + D^{\gamma}) \,, \tag{5}$$

where compute can be performed locally (C_L) or abroad in a remote data centre (C_R),

$$C = C_L + C_R. \tag{6}$$

We assume $0 < \alpha < 1$, $0 < \gamma < 1$, $D > 0$, and $C > 0$ so that compute and data are complements. We further assume diminishing marginal returns to data (i.e. $\gamma \in (0, 1]$), which is consistent with empirical (Bajari et al., 2019) and theoretical

(Farboodi and Veldkamp, 2021) economics, and motivated by power-law learning curves in machine learning (e.g. Hestness et al. (2017)). [7]

The restriction on cross-border data transfers enters the model through a compliance constraint:

$$C_L \geq \omega(\tau)C, \tag{7}$$

which simply states that at least a share $\omega(\tau)$ of compute needs to happen locally. $\frac{\partial \omega(\tau)}{\partial \tau} > 0$, and $0 < \omega(\tau) < 1$. We interpret tighter regulation as an increase in τ, which raises the required local compute share.

We assume the following cost structure. Regarding data, we abstract from the costs of generating and maintaining data-generating activity (such as product development and user acquisition). These are treated as fixed costs of opreating on the market, and are not affected by cross-border data transfer regulation.[8] At the same time, the firm faces costs for using compute. Remote compute is rented at a constant marginal cost p_R from a cloud service abroad.

$$\text{Cost}_R(C_R) = p_R C_R \tag{8}$$

Local compute however, has to be built (with a positive fixed cost F), and thus the average costs per unit of local compute decrease as C_L grows. Declining average costs of local compute arise exclusively from a fixed cost related to local compute. For simplicity, we abstract from technical economies of scale.

$$\text{Cost}_L(C_L) = F + p_L C_L. \tag{9}$$

Combining the production and costs functions yields the profit function:

$$\pi_c = N_c \left(C^\alpha \log(1 + D_c^\gamma)\right) - \text{Cost}_L(C_L) - \text{Cost}_R(C_R), \tag{10}$$

where N_c is the population in country c.

Proposition 1: Effective unit price under regulatory compliance. The impact of regulation on the unit price of compute depends on relative prices of local and remote compute. There are two options:

[7] Agrawal et al. (2022) note, that data can also have increasing marginal returns in firms' production function if it helps generate market power. Since firms are price-taking in our model, that channel omitted. We also note that the key comparative statics below do not rely on a particular assumption about the convexity or concavity of Y_c with respect to D.

[8] This simplified treatment mathces the standard view that, once collected, data are largely non-rival and cheap to reuse and the economically relevant margin is acquisition rather than per-use pricing Farboodi and Veldkamp (2021).

1. If $p_L \leq p_R$, firms' cost minimising allocation is always $C_L = C$, $C_R = 0$, and changes in τ do not affect firms' choices.

2. $p_L > p_R$, the constraint binds, and the effective unit price of compute can be expressed as
$$p(\omega) = p_R + (p_L - p_R)\omega(\tau), \tag{11}$$
and local and remote compute can be expressed as shares of total compute: $C_L = \omega(\tau)C$, $C_R = (1 - \omega(\tau))C$.[9]

Proof: In Appendix B.

According to Proposition 1, regulation only has an effect when $p_L > p_R$, or when remote computing is cheaper than local. This is the empirically relevant case for this paper, so throughout the rest of this section, we only concentrate on Case 2. In this case, regulation forces a share of the remote compute to be local, and, thus increases the effective unit price.

When the regulation bites, the optimal level of total compute $C^*(\tau)$ solves for the following first-order condition, where the marginal revenue equals marginal cost.
$$N_c \alpha (C^*)^{\alpha - 1} \log(1 + D_c^\gamma) = p_R + (p_L - p_R)\omega(\tau), \tag{12}$$
and solving for $C^*(\tau)$ gives the optimal level of compute.
$$C^*(\tau) = \left(\frac{N_c \, \alpha \, \log(1 + D_c^\gamma)}{p_R + (p_L - p_R)\, \omega(\tau)} \right)^{\frac{1}{1-\alpha}}. \tag{13}$$

Proposition 2. Impact of regulation on demand for total compute.
When $p_L > p_R$, tightening regulation decreases the total compute used, or $\dfrac{dC^*}{d\tau} < 0$.

Proof: Optimal C^* is given by Equation 13. Since $\omega'(\tau) > 0$, and $p_L > p_R$, then the denominator rises with τ, hence $C^*(\tau)$ falls.

Proposition 2 indicates, that, if regulation increases the price of compute, it also mechanically decreases its demand.

The effect of τ on the demand for local compute C_L, is ambiguous. Proposition 2 indicates, that regulation decreases the equilibrium amount of total compute.

[9] For simplicity, we assume, that all costs from regulation accrue to firms via their choice of local and remote compute. It would be possible to include compliance costs such as legal services and audits as an extra per-unit compliance cost term in the cost function. In that case, the effective unit costs could be expressed as $p(\omega(\tau)) = p_R + (p_L - p_R)\omega(\tau) + G(\tau)$, where $G'(\tau) > 0$. This shifts the threshold condition for when tighter regulation increases local compute, but it does not change the basic trade-off in the model. For conceptual clarity, we omit this channel.

But at the same time, the regulation pushes the share of local compute up. Which of these effects dominates, is given by Proposition 3.

When the regulation binds, the compliance constraint implies that optimal local compute is given by

$$C_L^*(\tau) = \omega(\tau)\, C^*(\tau). \tag{14}$$

Proposition 3. Impact of regulation on local compute.

When the compliance constraint binds, or $p_L > p_R$, $\dfrac{dC_L^*}{d\tau} > 0$ if and only if $p_R\,(1 - \alpha + \alpha\omega(\tau)) > \alpha\omega(\tau)p_L$.

Proof: In Appendix C.

Proposition 3 shows that when the price advantage of remote over local compute is sufficiently small, the increase in the required local share dominates the contraction in total compute demand, so tighter regulation raises local compute. Intuitively, regulation can shift the composition of demand toward local compute, which raises demand for domestic capacity even if it reduces the profit of compute users.[10]

Moreover, conditional on regulation increasing local compute, the magnitude of the response scales with market size and data intensity. This is formalised in Proposition 4. Proposition 4 demonstrates, that simple theoretical framework presented here can match the empirical patterns we observed.

Link to empirical measures. The model contains two country-level demand shifters, N_c and $\log(1 + D_c^{\gamma})$, that capture the heterogeneity patterns we observed empirically in Section 4.

First, N_c measures the market size. In Section 4, we proxy N_c using either population (market size in persons) or aggregate GDP (market size in USD, i.e. the average income per person multiplied by population size). Under either proxy, we found that the tightening of regulation a larger impact in C_L for larger market size.

Second, $\log(1 + D_c^{\gamma})$ captures the importance of data in the firm's revenues. In the empirical section, we proxied data intensity using an industry-level downstreamness measure. Higher downstreamness as reflects production that is closer to final demand and therefore more reliant on rich, granular customer transaction

[10] The model is partial equilibrium and isolates the intensive-margin response of local compute C_L holding the mass of firms fixed. When the compliance constraint binds, tighter regulation can reduce maximised profits, but we abstract from entry/exit decisions. We map the theory to the data by assuming that domestic data centre openings are monotonically increasing in aggregate local compute demand $\sum C_L = M\, C_L$ (with M fixed and tied to market size). Hence, when Proposition 3 implies higher C_L, it implies higher domestic capacity investment/openings.

data. In the model, this corresponds to a higher effective data term $\log(1 + D_c^\gamma)$, and thus larger level changes in local compute following a regulation shock.

Proposition 4. Level heterogeneity When the compliance constraint $p_L > p_R$ binds and the price level difference threshold $p_R\,(1 - \alpha + \alpha\omega(\tau)) > \alpha\omega(\tau)p_L$ holds, the effect of regulation on local compute can be expressed as

$$\frac{dC_L^*}{d\tau} = \omega'(\tau)C^*(\tau)\left[1 - \frac{\omega(\tau)}{1-\alpha}\frac{p_L - p_R}{p(\omega(\tau))}\right]. \tag{15}$$

The bracketed term only depends on prices, α and $\omega(\tau)$. Therefore, the magnitude of the response in C_L is increasing in $C^*(\tau)$, and increasing in market size N_c and in the data revenue shifter $\log(1 + D_c^\gamma)$.

In other words, for a common policy change $\Delta\tau$, countries with larger N_c, and larger $\log(1 + D_c^\gamma)$ experience larger level increases in local compute.

Proof: In Appendix C.

We have demonstrated, that the compliance constraint model can replicate the empirical patterns we observed in the data.

The model also delivers additional implications about *where* the compliance constraint bites the most. These implications arise from cross-country heterogeneity in the relative costs of local and remote compute. In the model, local energy and construction costs map into p_L, while international connectivity and access to cloud services map into p_R. We derive comparative statics with respect to (p_L, p_R) next and discuss how they translate into testable heterogeneity.

5.2 Additional Testable Implications

We formalise the testable implications in Proposition 5.

Proposition 5: Cost and connectivity heterogeneity

Assume $p_L > p_R$, so that the compliance constraint binds, and consider an interior solution.

$$B(\tau) \equiv 1 - \frac{\omega(\tau)}{1-\alpha}\frac{p_L - p_R}{p(\omega(\tau))} > 0,\ p(\omega(\tau)) \equiv p_R + (p_L - p_R)\omega(\tau).$$

In this case, the marginal effect of tightening on local compute is expressed as:

$$\frac{dC_L^*(\tau)}{d\tau} = \omega'(\tau)\, C^*(\tau)\, B(\tau).$$

1. Local cost monotonicity: Holding $(N_c, D_c, \alpha, \omega(\tau), \omega'(\tau))$ fixed, the marginal effect is strictly decreasing in p_L:

$$\frac{\partial}{\partial p_L}\left(\frac{dC_L^*(\tau)}{d\tau}\right) < 0. \tag{16}$$

   As a result, higher marginal cost of local compute (e.g. by higher electricity prices) results in a smaller reaction in C_L to tightened regulation.

2. Ambiguous impact of remote cost: Holding $(N_c, D_c, \alpha, \omega(\tau), \omega'(\tau))$ fixed, the effect is generally non-monotone in p_R because

$$\frac{\partial}{\partial p_R}\left(\frac{dC_L^*(\tau)}{d\tau}\right) = \omega'(\tau)\left[\underbrace{\frac{\partial C^*(\tau)}{\partial p_R}B(\tau)}_{<0} + \underbrace{C^*(\tau)\frac{\partial B(\tau)}{\partial p_R}}_{>0}\right]. \tag{17}$$

   An increase in p_R raises the effective unit price $p(\omega)$ and reduces C^* (the first term in brackets). At the same time, it also narrows the relative cost gap between remote and local compute (the second term). The relative magnitude of the two terms is unclear.

Proof: In Appendix D.

We explore these predictions empirically in the next section.

6 Empirical Test: Cost and Connectivity Heterogeneity

The theory model outlined above gives specific predictions about the heterogeneity related to cost-environment faced by firms in different countries. The model predicts that the marginal response of local compute is smaller when local compute is more expensive (higher p_L) and that the cost of remote compute (p_R) has an a priori ambiguous effect on local compute.

In this section, we take these predictions to the data. We proxy p_L with country-level electricity generation measured in 2010. The rationale is that data centres are electricity-intensive, so locations with greater generation capacity plausibly face fewer supply constraints and, on average, lower wholesale electricity costs for large industrial users. While generation is an imperfect proxy for prices (since it ignores fuel mix, regulation, and cross-border power trade), it captures a basic constraint on the scale of energy supply that is relevant for hosting electricity-intensive facilities. At the same time, electricity generation is highly correlated with GDP and population (Table 4), so the split may partially capture market-size heterogeneity already documented earlier.

We measure the cost of remote compute, p_R, using international internet connectivity. For this, we use the international bandwidth usage measured by International Telecommunications union[11]. The rationale is that using compute abroad requires reliable, high-capacity international links: greater bandwidth reduces effective frictions in accessing remote services (e.g., congestion and quality constraints) and is correlated with better integration into global digital networks. As with electricity generation, bandwidth is an imperfect proxy (it does not directly measure cloud prices or latency), but it captures an observable component of the infrastructure needed to rely on remote compute at scale. International bandwidth is likewise strongly related to income and market size (Table 4), implying that the "high bandwidth" group overlaps substantially with richer economies.

As before, to alleviate concerns related to reverse causality, we fix these measures at their 2010 values. Moreover, due to small sample sizes, we dichotomise the variables using a simple below/above median split. Because these splits are correlated with GDP and population, the event studies below should be interpreted as comparisons of higher- vs lower-electricity (or bandwidth) environments that may also differ in market size; they provide a check of whether the model's comparative statics appear in the data rather than a decomposition of mechanisms.

The resulting event studies are given in Figure 7. In panel (a), we report, that tightening of data regulation has, on average, a larger effect on new data centre buildout in countries wiht high electricity production (and presumably, lower energy prices). This is consistent with theory outlined above.

For the remote-connectivity split (Figure 7 (b)), we find that data centre buildout after tightening is concentrated in countries with high international bandwidth. In the model, this is consistent with a domestic substitution mechanism: when cross-border connectivity is good, remote compute is effectively cheaper and firms rely more on it; tightening then requires a larger share of activity to be executed domestically, raising C_L in the tightening countries.

In terms of Proposition 5, higher connectivity lowers the effectve marginal cost of remote compute p_R (where p_R is a combination of cloud rental price plus a connectivity or latency component). In the expression, $\frac{dC_L^*(\tau)}{d\tau} = \omega'(\tau)\, C^*(\tau)\, B(\tau)$, a lower p_R raises the term C^* by lowering the effective unit price, but also widens the local-remote price gap $(p_L - p_R)$ which pushes down the final term $B(\tau)$.

[11]The ata are downloaded from https://datahub.itu.int/data/?e=MLI&i=242

The observation that high-connectivity countries react more, implies, that, in the region relevant in our data remote and local compute costs are sufficiently close, so tightening mainly shifts compute from abroad to domestic sites without a large contraction in total compute demand within the tightening country.

In Figure 7 (b), the low-connectivity group shows post-treatment coefficients that are at or slightly below zero. Interpreted literally, this says that countries with weak international connectivity, stricter regulation does not result in domestic buildout. On the contrary, it leads to a decline in the flow of new domestic data centres.

In the language of the theory model, this pattern fits the region where $p_L > p_R$, and remote compute is not a meaningful margin to reallocate from. In that region, tightening would theoretically predict a zero effect on buildout. A systematically negative response might be seen as an additional friction that shrinks the production of data-intensive goods rather than triggering a substitution into domestic compute.

7 Discussion

Taken together, the theoretical compliance constraint model and our empirical estimates imply that the tightening of data transfer laws results in fragmentation of compute. There are more data centre sites built as a result of the tightening of regulation in the tightening countries-

This also implies a loss of real resources. The fragmentation results in parallel "stacks" across jurisdictions: duplicated investments, duplicated staffing, duplicated network buildout, and duplicated compliance processes. Even if total compute use is likely to decrease as a result of stricter regulation, average costs increase because fixed costs are spread over a smaller scale.

Using data transfer restrictions as industrial policy to attract domestic data centre investment is therefore costly for local firms. In the model, there are two cases. If local compute is cheaper than remote compute ($p_L \leq p_R$), the compliance constraint is slack and regulation has little effect on firms' allocation choices or investment outcomes, because firms would already choose to process locally. On the other hand, if remote compute is cheaper ($p_L > p_R$), the constraint binds and firms are forced to shift part of their workload onshore. This induces two cost increases. Firms incur the fixed cost of maintaining compliant local capacity (or contracting for it), which can trigger additional local site entry. Additionally, the effective marginal cost of compute rises because a larger share of compute must be

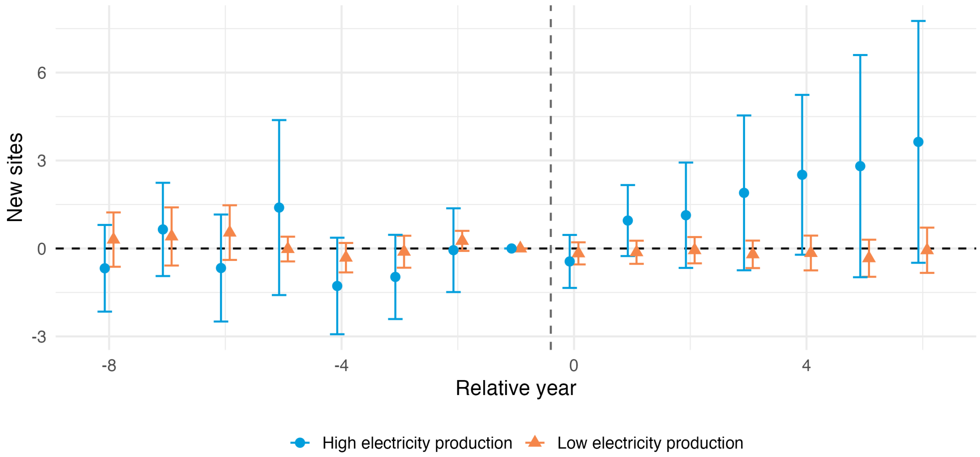

(a) Event study: electricity price split

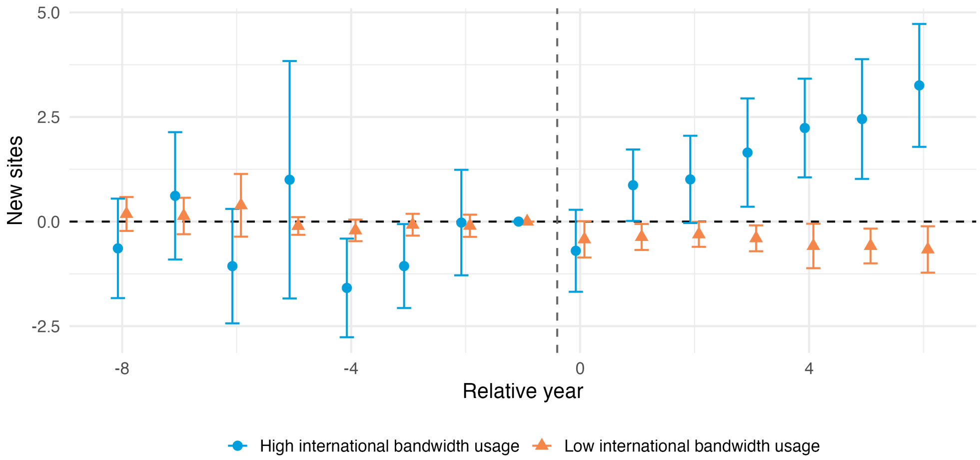

(b) Event study: international bandwidth split

Figure 7. Event study pictures. The vertical lines correspond to 95% confidence intervals.

purchased at the higher local price. In this sense, data transfer restrictions function like a tax on data intensive production: they can raise domestic investment in infrastructure, but also increase the unit cost of compute for local firms.

With fixed costs, the choice to host locally is easier to accommodate when local demand is large. The compliance constraint generates a size gradient with regard to local market size and local energy prices. Larger markets can host local "hyperscale-like" facilities, while smaller markets have to rely on smaller, higher-cost data centres. Regulation can result in some small or high-cost countries to be effectively excluded. As a result of regulation, a small country with high fixed costs might end up in a "corner solution" with no meaningful capacity.

To understand the quantitative implications for the treated countries, we build a simple back-of-the-envelope counterfactual based on the cohort-averaged ATT estimate Table 5 Column (1)). For each treated country-year, we construct a counterfactual number of new sites by subtracting the estimated treatment effect from the observed new sites, and then cumulate this number through 2024. The results of this exercise, plotted, in Figure 8 suggest that, absent the tightening of data flow restrictions, the number of data centre sites in treated countries would be appoximately 15% lower (relative to the number of data centre sites in the treated countries). Moreover, this aggregate figure masks substantial heterogeneity across countries. The impact is considerably more pronounced in larger and wealthier countries with larger consumer-facing markets.[12]

Despite the strong assumptions embedded in this counterfactual, the magnitude still suggests that data-flow restrictions have had a noticeable impact on new data centre buildout. Consistent with this interpretation, our regional spillover tests suggest no reallocation to neighbouring countries.

It is worth contrasting our work to previous research studying how data regulation affects firms' data collection and use Demirer et al. (2024); Chang et al. (2023). Those papers emphasise a "wedge" channel: higher compliance costs lead firms to use and retain less data and — via complementarity — to reduce compute used for data intensive activity. Our model can reproduce this reduction in total compute, but it highlights an additional margin of adjustment: regulation can also reallocate compute geographically by shifting *where* compute is hosted. The wedge channel operates primarily through the level of activity, while the compliance channel operates through the composition of compute across locations; these

[12] This counterfactual is a mechanical accounting exercise anchored on site counts (not capacity, output, or welfare), and it abstracts from general-equilibrium relocation of investment across countries.

forces can move domestic infrastructure investment in opposite directions.

As with any study relying on country-level data, we are faced with issues related to small sample sizes. The number countries in the world — let alone countries that have changed their data transfer legislation — is naturally limited. The small sample sizes combined with a highly right-skewed dependent variable and potential measurement error in the classification of the data transfer legislation all work against precision when trying to find statistically significant effects. Nonetheless, the main patterns are relatively robust to a set of reasonable sensitivity checks. In particular, the fact that our results remain qualitatively similar using winsorised data implies that our results are not only driven by a small number of extreme observations.

Another potential concern related to small samples is the informational content of the heterogeneity splits. The heterogeneity analyses underlying the theory always rely on splitting the data into high/low groups by the group median. A critical reader could have a valid concern that the splits produce highly overlapping country groups, limiting how cleanly each split can be interpreted. This could create issues analogous to multicollinearity: when the groupings contain the same countries, it becomes hard to attribute differences in estimated effects to any one of the overlapping dimensions. We, however, find that the correlation between most of the splits is, at most, moderate, demonstrating that while correlated, the different splits are not simply capturing the same underlying country grouping. Therefore, the heterogeneity patterns should be read as evidence consistent with cost and scale based comparative statistics, rather than cleanly separable causal channels.

Our empirical and theoretical results rely on a set of simplifying assumptions that stem from limitations related to underlying data. First, our outcome is the number of sites, which does not exactly map to aggregate data centre capacity in any country. Therefore, an increase in sites should not be mechanically interpreted as an increase in total capacity; it is best interpreted as evidence of increased entry and a higher likelihood of duplicative fixed costs.

Second, we do not differentiate between the types of compute available locally and remotely. For instance, frontier AI training requires specialised hardware such as advanced GPU chips, access to which can be geographically concentrated and, in some jurisdictions, constrained by export controls (U.S. Department of Commerce, Bureau of Industry and Security, 2022). When these non-modelled constraints bind, a compliance quota like the one modelled in this paper can increase domestic sites and shift compliance-sensitive workloads onshore, while leaving reliance on foreign frontier compute largely unchanged. While relevant in practice, we abstract away from these considerations.

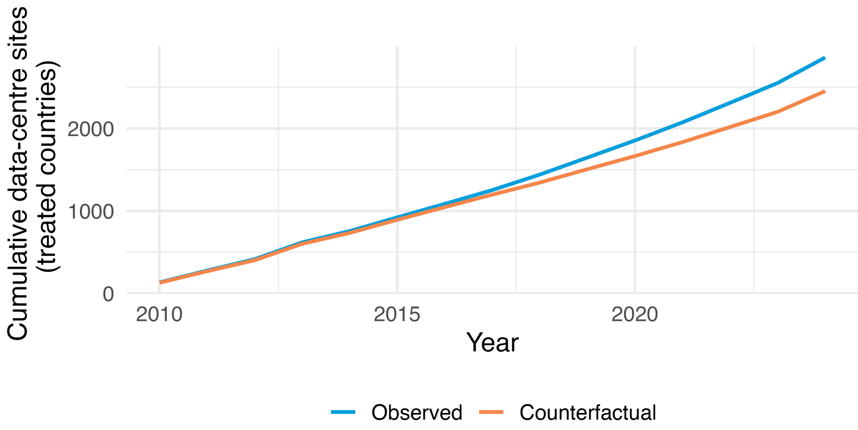

Figure 8. Cumulative data-centre sites: observed versus counterfactual without tightening of cross-border data-transfer restrictions.

8 Conclusions

In this paper, we examine how restrictions on cross-border data transfers affect the location of data centre investment. Using a Sun–Abraham two-way fixed effects event-study design (Sun and Abraham, 2021) on a global country–year panel of data centre openings and law changes, we find that tightenings of data transfer rules are followed by higher domestic data centre entry in the tightening countries. The dynamic estimates show stable pre-trends, and a within-region leave-one-out placebo test provides no evidence of systematic changes in nearby countries' entry around the tightening date, supporting the interpretation of the estimates as a domestic entry response.

To rationalise these findings, we present a simple micro-foundation model where regulation affects firms' choices through a compliance quota, which states that at least a certain share of compute needs to take place locally. We demonstrate that this model can replicate a wide range of comparative statics we estimate from the data.

Using our CATT estimates, we find that, in the absence of regulation, the number of data centres in the countries that have implemented regulation would have been approximately 15% lower. Our results highlight that regulation has real economic costs, and that a smart regulator should account for potential costs of regulation already ex ante.

Finally, we highlight, that despite the these costs, laws restricting data transfers can still be rationalised as insurance against surveillance or coercion by foreign hubs. This maps to the "weaponised interdependence" argument: countries may accept higher costs to reduce exposure to foreign-controlled network hubs (Farrell and Newman, 2019). Our evidence provides systematic support for the premise that such rules are indeed followed by domestic buildout.

References

Agrawal, A., Gans, J. and Goldfarb, A. (2019), *The economics of artificial intelligence: An agenda*, University of Chicago Press.

Agrawal, A., Gans, J. and Goldfarb, A. (2022), *Prediction machines, updated and expanded: The simple economics of artificial intelligence*, Harvard Business Press.

Amiti, M., Kong, S. H. and Weinstein, D. (2020), The effect of the us-china trade war on us investment, Technical report, National Bureau of Economic Research.

Antràs, P., Chor, D., Fally, T. and Hillberry, R. (2012), 'Measuring the upstreamness of production and trade flows', *American Economic Review* **102**(3), 412–416.

Arel-Bundock, V., Enevoldsen, N. and Yetman, C. (2018), 'countrycode: An r package to convert country names and country codes', *Journal of Open Source Software* **3**(28), 848.
URL: *https://doi.org/10.21105/joss.00848*

Aridor, G., Che, Y.-K. and Salz, T. (2020), *The economic consequences of data privacy regulation: Empirical evidence from GDPR*, National Bureau of Economic Research Cambridge, MA, USA.

Bajari, P., Chernozhukov, V., Hortaçsu, A. and Suzuki, J. (2019), The impact of big data on firm performance: An empirical investigation, *in* 'AEA papers and proceedings', Vol. 109, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 33–37.

Bonfiglioli, A., Crinò, R., Filomena, M. and Gancia, G. (2025), Data, power and emissions: The environmental cost of ai, Technical report, CESifo Working Paper.

Caldara, D., Iacoviello, M., Molligo, P., Prestipino, A. and Raffo, A. (2020), 'The economic effects of trade policy uncertainty', *Journal of Monetary Economics* **109**, 38–59.

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

Casalini, F. and López González, J. (2019), Trade and cross-border data flows, OECD Trade Policy Papers 220, OECD Publishing, Paris.
URL: *http://dx.doi.org/10.1787/b2023a47-en*

Chang, Q., Cong, L. W., Wang, L. and Zhang, L. (2023), 'Production, trade, and cross-border data flows', *Available at SSRN 4672185* .

Chari, V. V., Kehoe, P. J. and McGrattan, E. R. (2007), 'Business cycle accounting', *Econometrica* **75**(3), 781–836.

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

Demirer, M., Hernández, D. J. J., Li, D. and Peng, S. (2024), Data, privacy laws and firm production: Evidence from the gdpr, Technical report, National Bureau of Economic Research.

Di Tella, R. and Schargrodsky, E. (2004), 'Do police reduce crime? estimates using the allocation of police forces after a terrorist attack', *American Economic Review* **94**(1), 115–133.

European Parliament and Council of the European Union (2016), 'Regulation (eu) 2016/679 (general data protection regulation)', Official Journal of the European Union, L 119, 4.5.2016, pp. 1–88.
URL: *https://eur-lex.europa.eu/eli/reg/2016/679/oj*

Farboodi, M. and Veldkamp, L. (2021), A model of the data economy, Technical report, National Bureau of Economic Research Cambridge, MA, USA.

Farrell, H. and Newman, A. L. (2019), 'Weaponized interdependence: How global economic networks shape state coercion', *International security* **44**(1), 42–79.

Ferracane, M. F. (2017), Restrictions on cross-border data flows, Technical report, ECIPE working paper.

Frey, C. B. and Presidente, G. (2024), 'Privacy regulation and firm performance: Estimating the gdpr effect globally', *Economic Inquiry* **62**(3), 1074–1089.

Goldberg, S. G., Johnson, G. A. and Shriver, S. K. (2024), 'Regulating privacy online: An economic evaluation of the gdpr', *American Economic Journal: Economic Policy* **16**(1), 325–358.

Goldfarb, A. and Tucker, C. (2019), 'Digital economics', *Journal of economic literature* **57**(1), 3–43.

Greenleaf, G. (2021), Global data privacy laws 2021: Despite covid delays, 145 laws show gdpr dominance, Technical Report UNSW Law Research Paper No. 21-60, University of New South Wales, Faculty of Law. (2021) 169 Privacy Laws & Business International Report, 1, 3–5. Published February 11, 2021.
URL: *https://ssrn.com/abstract=3836348*

Greenleaf, G. (2025), Global data privacy laws 2025: 172 countries, twelve new in 2023/24, Technical report, University of New South Wales, Faculty of Law. Published April 2, 2025.
URL: *https://ssrn.com/abstract=5275559*

Gupta, S., Ghosh, P. and Sridhar, V. (2022), 'Impact of data trade restrictions on it services export: A cross-country analysis', *Telecommunications Policy* **46**(9), 102403.

Handley, K. and Limao, N. (2015), 'Trade and investment under policy uncertainty: theory and firm evidence', *American Economic Journal: Economic Policy* **7**(4), 189–222.

Handley, K. and Limão, N. (2017), 'Policy uncertainty, trade, and welfare: Theory and evidence for china and the united states', *American economic review* **107**(9), 2731–2783.

Hestness, J., Narang, S., Ardalani, N., Diamos, G., Jun, H., Kianinejad, H., Patwary, M. M. A., Yang, Y. and Zhou, Y. (2017), 'Deep learning scaling is predictable, empirically', *arXiv preprint arXiv:1712.00409* .

Hsieh, C.-T. and Klenow, P. J. (2009), 'Misallocation and manufacturing tfp in china and india', *The Quarterly journal of economics* **124**(4), 1403–1448.

Janßen, R., Kesler, R., Kummer, M. E. and Waldfogel, J. (2022), Gdpr and the lost generation of innovative apps, Technical report, National Bureau of Economic Research.

Jardim, E., Long, M. C., Plotnick, R., Vigdor, J. and Wiles, E. (2024), 'Local minimum wage laws, boundary discontinuity methods, and policy spillovers', *Journal of Public Economics* **234**, 105131.

Jia, J., Jin, G. Z., Leccese, M. and Wagman, L. (2025), How does privacy regulation affect transatlantic venture investment? evidence from gdpr, Technical report, National Bureau of Economic Research.

Johnson, G. A., Shriver, S. K. and Goldberg, S. G. (2023), 'Privacy and market concentration: Intended and unintended consequences of the gdpr', *Management Science* **69**(10), 5695–5721.

Koski, H. and Valmari, N. (2020), Short-term impacts of the gdpr on firm performance, Technical report, ETLA Working Papers.

Kwon, C.-W. and Chun, B. G. (2009), 'Local content requirement under vertical technology diffusion', *Review of Development Economics* **13**(1), 111–124.

Lehdonvirta, V., Wu, B. and Hawkins, Z. (2025), 'Weaponised interdependence in a bipolar world: how economic forces and security interests shape the global reach of us and chinese cloud data centres', *Review of International Political Economy* pp. 1–26.

Ling, D. C., Naranjo, A. and Scheick, B. (2021), 'There is no place like home: Information asymmetries, local asset concentration, and portfolio returns', *Real Estate Economics* **49**(1), 36–74.

Ling, D. C., Wang, C. and Zhou, T. (2021), 'The geography of real property information and investment: Firm location, asset location and institutional ownership', *Real Estate Economics* **49**(1), 287–331.

Mancini, M., Montalbano, P., Nenci, S. and Vurchio, D. (2024), 'Positioning in global value chains: World map and indicators, a new dataset available for gvc analyses', *The World Bank Economic Review* **38**(4), 669–690.

Pan Fang, T. and Greenstein, S. (2025), 'Where the cloud rests: The economic geography of data centers', *Strategy Science* **10**(4), 404–420.

Peukert, C., Bechtold, S., Batikas, M. and Kretschmer, T. (2022), 'Regulatory spillovers and data governance: Evidence from the gdpr', *Marketing Science* **41**(4), 746–768.

Pierce, J. R. and Schott, P. K. (2018), 'Investment responses to trade liberalization: Evidence from us industries and establishments', *Journal of International Economics* **115**, 203–222.

Restuccia, D. and Rogerson, R. (2008), 'Policy distortions and aggregate productivity with heterogeneous establishments', *Review of Economic dynamics* **11**(4), 707–720.

Sisto, E. and Van der Marel, E. (2025), Privacy at a price? an empirical analysis of gdpr's impact on eu trade flows, Technical report, ECIPE Occasional Paper.

Sun, L. and Abraham, S. (2021), 'Estimating dynamic treatment effects in event studies with heterogeneous treatment effects', *Journal of econometrics* **225**(2), 175–199.

Sun, R. and Trefler, D. (2023), The impact of ai and cross-border data regulation on international trade in digital services: A large language model, Technical report, National Bureau of Economic Research.

Tian, C., Chellappa, R. K. and Martinez, J. G. (2025), Supply, demand or policy driven? an empirical examination of data center location strategies in the united states, *in* 'Proceedings of the International Conference on Information Systems (ICIS 2025)', Association for Information Systems (AIS).
URL: *https://aisel.aisnet.org/icis2025/public˙is/public˙is/10*

U.S. Department of Commerce, Bureau of Industry and Security (2022), 'Bis updated public information page on export controls imposed on advanced computing and semiconductor manufacturing items to the people's republic of china (prc)'. Accessed January 7, 2026.
URL: *https://www.bis.gov/press-release/bis-updated-public-information-page-export-controls-imposed-advanced-computing-semiconductor*

Zhang, H., Kim, H. and Moon, S. (2025), 'The effects of data restriction policies and institutions on digital financial service and trade', *Emerging Markets Finance and Trade* **61**(9), 2745–2762.

A Appendix: Restrictiveness Classification

The following stepwise procedure was applied to determine each country's OECD restrictiveness level for cross-border data transfers.

The logic proceeds sequentially from basic legal coverage to the specific transfer conditions, adequacy recognition, fallback safeguards, and authorization requirements.

Stage 1 – Legal Coverage (Level 0)

1. Check whether the country's law regulates **personal data** and includes provisions on **cross-border data transfers**.
2. If *no*, assign **Level 0 (No regulation)**.
3. If *yes*, proceed to Stage 2.

Stage 2 – Type of Legal Condition

1. Determine whether cross-border transfers are:
   - **Free** → **Level 1 (Ex-post accountability)**.
   - **Conditional** (subject to adequacy or safeguards) → Stage 3.
   - **Authorized** (require prior regulator approval) → Stage 4.

Stage 3 – Conditional Transfers (Levels 2–4) For conditional regimes, classification depends on who determines adequacy, whether an adequacy list exists, and what fallback mechanisms are recognized.

1. Identify **who determines adequacy:**
   - **Private assessment** (by data exporter) → **Level 2 (Self-assessment + fallbacks)**.
   - **Public assessment** (by data protection authority) → Continue to Step 2.
2. Verify whether the law establishes or references an official **list of adequate jurisdictions.**
   - If such a list is maintained by the data protection authority or government, this confirms a **public adequacy mechanism** (Levels 3–4).
3. Verify the presence of **fallback mechanisms** available when adequacy is not established, such as:
   - Consent of the data subject,
   - Standard Contractual Clauses (SCCs),
   - Binding Corporate Rules (BCRs),

- Approved certifications or codes of conduct,
- Public interest or necessity exceptions.

4. Classify the level of conditional restrictiveness:
   - If fallback mechanisms are legally recognized → **Level 3 (Public assessment + options)**.
   - If the law further imposes downstream processing or localization obligations (e.g., requiring that foreign recipients apply the sender's domestic rules or holding exporters liable for foreign processors) → **Level 4 (Public assessment + processing conditions)**.

Stage 4 – Authorization Regimes (Levels 5–6) For authorization-based regimes, the level depends on the scope of regulatory approval required and the presence (or absence) of adequacy or fallback pathways.

1. Determine whether prior authorization applies:
   - Only for **non-adequate destinations**, while adequate jurisdictions or fallback safeguards (e.g., SCCs, BCRs, consent) remain available → **Level 5 (Case-by-case authorization + adequacy fallback)**.
   - For **all transfers** or if **local storage/processing** of data is mandatory, with no fallback options permitted → **Level 6 (Universal authorization / localization)**.

Outcome: The final OECD level corresponds to the **maximum restrictiveness** observed for general personal data transfers in each jurisdiction. This framework integrates both the **primary transfer conditions** and the **availability of fallback mechanisms** to ensure consistent, rule-based classification.

B Proof of Proposition 1

Proof. Fix total compute $C > 0$ and consider the cost-minimisation problem over the allocation (C_L, C_R):

$$\min_{C_L, C_R \geq 0} \ \text{Cost}_L(C_L) + \text{Cost}_R(C_R) \quad \text{s.t.} \quad C_L + C_R = C, \quad C_L \geq \omega(\tau)C,$$

where $\text{Cost}_L(C_L) = F + p_L C_L$ and $\text{Cost}_R(C_R) = p_R C_R$.

Using the accounting identity $C_R = C - C_L$, the objective can be written as

$$F + p_L C_L + p_R(C - C_L) = F + p_R C + (p_L - p_R)C_L.$$

Since $F + p_R C$ is constant for fixed C, the problem reduces to choosing C_L to minimise $(p_L - p_R)C_L$ subject to the feasibility interval $C_L \in [\omega(\tau)C,\ C]$.

Case 1: $p_L \leq p_R$. Then $p_L - p_R \leq 0$, so the objective is weakly decreasing in C_L. The cost-minimising choice is $C_L = C$ and $C_R = 0$. The compliance constraint is slack, so changes in τ (and thus $\omega(\tau)$) do not affect the allocation.

Case 2: $p_L > p_R$. Then $p_L - p_R > 0$, so the objective is strictly increasing in C_L. The cost-minimising choice is the smallest feasible C_L, i.e. the compliance constraint binds:

$$C_L = \omega(\tau)C, \qquad C_R = (1 - \omega(\tau))C.$$

Substituting this allocation into variable cost gives

$$p_L C_L + p_R C_R = p_L \omega(\tau)C + p_R(1 - \omega(\tau))C = [p_R + (p_L - p_R)\omega(\tau)]\, C.$$

Defining $p(\omega) = p_R + (p_L - p_R)\omega(\tau)$ gives the effective unit price $p(\omega)C$. □

C Proof of Proposition 3

Proof. Assume that the regulation binds, i.e. $p_L > p_R$, so that $C_L = \omega(\tau)C$ and $C_R = (1 - \omega(\tau))C$.

Step 1: Reduced-form profit and optimal C^*.

Since $C = C_L + C_R$, profits can be expressed as

$$\begin{aligned}\pi(C;\tau) &= N_c\, C^\alpha \log(1 + D_c^\gamma) - F - p_L\, \omega(\tau)\, C - p_R\, (1 - \omega(\tau))\, C \\ &= N_c\, C^\alpha \log(1 + D_c^\gamma) - F - p(\omega(\tau))\, C,\end{aligned}$$

where $p(\omega) = p_R + (p_L - p_R)\omega$.

The F.O.C for an interior optimum C^* is

$$N_c\, \alpha\, (C^*)^{\alpha-1} \log(1 + D_c^\gamma) = p(\omega). \tag{C.1}$$

This implies, that

$$C^*(\omega) = \left(\frac{N_c\, \alpha\, \log(1 + D_c^\gamma)}{p(\omega)}\right)^{\frac{1}{1-\alpha}}. \tag{C.2}$$

Step 2: Local compute and its derivative with respect to ω. In the binding regime,

$$C_L^*(\omega) = \omega\, C^*(\omega). \tag{C.3}$$

Differentiate with respect to ω:

$$\frac{dC_L^*}{d\omega} = C^*(\omega) + \omega\, \frac{dC^*}{d\omega}. \tag{C.4}$$

From (C.2), using $p'(\omega) = p_L - p_R$,

$$\frac{dC^*}{d\omega} = -\frac{1}{1-\alpha}\, C^*(\omega)\, \frac{p'(\omega)}{p(\omega)} = -\frac{1}{1-\alpha}\, C^*(\omega)\, \frac{p_L - p_R}{p(\omega)}.$$

Substitute into (C.4):

$$\frac{dC_L^*}{d\omega} = C^*(\omega) \left[1 - \frac{\omega}{1-\alpha} \frac{p_L - p_R}{p(\omega)}\right]. \tag{C.5}$$

Step 3: Obtain the threshold for the partial derivative.

Since $C^* > 0$, and $\omega'(\tau) > 0$, the sign of $\dfrac{dC_L^*}{d\omega}$ determined by the sign of the bracketed term.

C_L increases with ω (and also with τ), if:

$$1 - \frac{\omega}{1-\alpha} \frac{p_L - p_R}{p(\omega)} > 0. \tag{C.6}$$

substituting $p(\omega) = p_R + (p_L - p_R)\omega$ and rearranging gives:

$$p_R\,(1 - \alpha + \alpha\omega(\tau)) > \alpha\omega(\tau)p_L, \tag{C.7}$$

which is the desired expression. □

D Proof of Proposition 5

Proof. Fix τ so that $\omega = \omega(\tau) \in (0,1)$ and $\omega' = \omega'(\tau) > 0$ are constants, and assume $p_L > p_R$ so the compliance constraint binds. Then $C_L = \omega C$, $C_R = (1-\omega)C$, and the effective unit price of compute is

$$p(\omega) = p_R + (p_L - p_R)\omega = \omega p_L + (1-\omega)p_R.$$

Let $A = \frac{\omega}{1-\alpha} > 0$ and define

$$B(\tau) = 1 - A\frac{p_L - p_R}{p(\omega)}.$$

From Proposition 4 for an interior solution,

$$\frac{dC_L^*}{d\tau} = \omega' \, C^*(\tau) \, B(\tau).$$

(i) Local cost p_L monotonicity. Since

$$C^*(\tau) = \left(\frac{\alpha N_c \log(1 + D_c^\gamma)}{p(\omega)}\right)^{\frac{1}{1-\alpha}}, \qquad \frac{\partial C^*}{\partial p(\omega)} = -\frac{1}{1-\alpha}\frac{C^*}{p(\omega)} < 0,$$

and $\frac{\partial p(\omega)}{\partial p_L} = \omega$, we have

$$\frac{\partial C^*}{\partial p_L} = -\frac{\omega}{1-\alpha}\frac{C^*}{p(\omega)} < 0.$$

Moreover,

$$\frac{\partial B}{\partial p_L} = -A\frac{p(\omega) - (p_L - p_R)\,\partial p(\omega)/\partial p_L}{p(\omega)^2} = -A\frac{p(\omega) - (p_L - p_R)\omega}{p(\omega)^2} = -\frac{\omega}{1-\alpha}\frac{p_R}{p(\omega)^2} < 0.$$

Therefore,

$$\frac{\partial}{\partial p_L}\left(\frac{dC_L^*}{d\tau}\right) = \omega'\left(\frac{\partial C^*}{\partial p_L}B + C^*\frac{\partial B}{\partial p_L}\right) = -\omega' C^* \frac{\omega}{1-\alpha}\left(\frac{B}{p(\omega)} + \frac{p_R}{p(\omega)^2}\right).$$

Under interior condition $B(\tau) > 0$ (equivalently $p_R(1 - \alpha + \alpha\omega) > \alpha\omega p_L$), the bracket is strictly positive, hence

$$\frac{\partial}{\partial p_L}\left(\frac{dC_L^*}{d\tau}\right) < 0.$$

(ii) Remote cost p_R is ambiguous. We have $\frac{\partial p(\omega)}{\partial p_R} = 1 - \omega > 0$, so

$$\frac{\partial C^*}{\partial p_R} = \frac{\partial C^*}{\partial p(\omega)}\frac{\partial p(\omega)}{\partial p_R} = -\frac{1-\omega}{1-\alpha}\frac{C^*}{p(\omega)} < 0.$$

Also,

$$\frac{\partial}{\partial p_R}\left(\frac{p_L - p_R}{p(\omega)}\right) = \frac{-p(\omega) - (p_L - p_R)(1-\omega)}{p(\omega)^2} = -\frac{p_L}{p(\omega)^2} < 0,$$

so

$$\frac{\partial B}{\partial p_R} = -A\frac{\partial}{\partial p_R}\left(\frac{p_L - p_R}{p(\omega)}\right) = \frac{\omega}{1-\alpha}\frac{p_L}{p(\omega)^2} > 0.$$

Thus

$$\frac{\partial}{\partial p_R}\left(\frac{dC_L^*}{d\tau}\right) = \omega' \left[\underbrace{\frac{\partial C^*}{\partial p_R}}_{<0} B + \underbrace{C^* \frac{\partial B}{\partial p_R}}_{>0}\right],$$

which is of ambiguous sign. □