Firm in the clouds

Supplier market power and technology adoption with evidence from cloud computing

Nicholas Emery-Xu[‡]

October 19, 2026

Updated frequently; click here for the latest version

Abstract

I study the impact of supplier concentration on downstream technology adoption and firm productivity in the rapidly expanding French cloud computing services industry. Supplier concentration can reduce adoption and raise costs on adopting firms, slowing productivity growth. In addition, market power allows suppliers to levy egress fees on data moved off-platform or introduce additional complementarity between their cloud platforms and on-premise software, foreclosing rivals. Estimating a dynamic model of firm production and supplier pricing to study how downstream firms make technology bundling choices, I analyze the effects of competition policies on adoption and welfare. I find that adopting cloud computing increases firm productivity by 0.2% to 1.6%, with heterogeneous effects across sectors and a larger impact for firms that purchase services from multiple cloud providers. Entry costs are rapidly falling, while switching costs are around €3/GB of data stored. Simulating the impacts of a new entrant at the lowest quality level, I find welfare rises by 2.82% annually (€8.4 billion), while a merger of Google's and Microsoft's cloud platforms lowers welfare by 0.29% (€905 million). Simulating the effects of a ban on data egress fees from the EU Data Act and an increase in software interoperability between Microsoft server software and rival cloud platforms, I find these policies produce annual welfare gains of 0.06% (€193 million) and 0.16% (€517 million).

[*]UCLA, Department of Economics

[†]I am grateful for the support and feedback of my advisors John Asker, Hugo Hopenhayn, Simon Board, and Will Rafey. I thank Michael Rubens, Jackson Mejia, Danial Lashkari, Martin Hackmann, Yangkeun Yun, Bill Whyman, Jiaer Liu, Theodore Naff, Jacopo Bellabio, Jihye Jeon, Jacob Lariviere, Ilia Murtazashvili and seminar and workshop participants at UCLA, MIT FutureTech, IIOC 2026, BSE Summer Forum Applied IO Workshop 2026, and the Institute for Humane Studies for data provision and useful feedback. I thank several anonymous cloud engineers at a large health insurance firm, and a data scientist at a cloud security firm for answering my questions about adoption. I thank the Institute for Humane Studies, UCLA Lewis L Clarke Fellowship, and Coefficient Giving for financial support. This work used computational and storage services associated with the Hoffman2 Cluster which is operated by the UCLA Office of Advanced Research Computing's Research Technology Group. All remaining errors are my own.

1 Introduction

> *"By 2030: 75% of European businesses should use cloud-edge technologies for their activities."* - European Commission

> *"The French position is to encourage and support European innovation while preventing historical monopolies from recurring."* - Mireille Clapot, MP

General-purpose technologies increase firm productivity, but may be adopted too slowly as firms continue to use older technologies (Collard-Wexler and De Loecker (2015); Chari and Hopenhayn (1991); Bresnahan and Trajtenberg (1995)). Existing literature largely focuses on adoption frictions (Hornbeck et al. (2024)), assuming technology input decisions are independent of supplier market power.

However, many information technologies are provided by a relatively small handful of suppliers, raising concerns that market power could raise input markups or price discriminate against users, producing inefficiently low rates of adoption of these technologies. A recent example is cloud computing, which allows firms to increase productivity by leasing computing inputs over the Internet (Jin and Bai (2022)). The general-purpose nature of the technology, combined with a high concentration of suppliers, has prompted regulatory agencies worldwide to begin to research how to increase supplier competition with the goal of increasing utilization (Competition and Markets Authority (2025h); Korea Fair Trade Commission (2022); Japan Fair Trade Commission (2022)).

In this paper, I ask whether supplier market power produces inefficiencies in the rate of technology adoption in the cloud computing industry. To do so, I construct a new model of firm technology adoption in which firms choose whether to purchase cloud computing services from a small set of differentiated suppliers. Adopting cloud computing offers productivity benefits, which are compounded when firms bundle suppliers. Firms face two consequences of market power, prices above marginal cost (Byrne et al. (2018)) and switching frictions when changing suppliers, which, which partly technological, are also driven by suppliers' data transfer fees. I derive a computationally-tractable estimation strategy to structurally estimate the model parameters. Using the model, I assess the counterfactual impact of competition policies proposed by regulators on downstream adoption and total welfare.

I introduce supplier choice in a model of industry dynamics with technology adoption. Each year, firms choose their set of computing inputs, as well as capital investment decisions in anticipation of their future productivity gains. Firms can choose between on-premise only computing (old technology) or on-premise along with any bundle of cloud technologies from existing providers. Adopting cloud computing increases firm productivity, with addi-

tional benefits from bundling multiple providers, and providers are heterogeneous in quality-adjusted variable costs. Finally, firms face fixed costs when first adopting cloud computing and when switching bundles that serve as a form of lock-in to existing technologies.

To model supplier market power, vertically-differentiated cloud computing suppliers compete for short run profits by setting per-unit computing prices and charging data egress fees for moving data between suppliers. After observing the prior period market shares, firms choose egress fees and usage prices to attract new consumers and maximize revenue from existing ones.

The modeled effects that can produce inefficiently low adoption rates are twofold. First, bundling cloud providers creates positive productivity benefits for users that are at best partially priced in by incumbent firms (Bresnahan and Trajtenberg (1995)). In addition, the high concentration on the supply side raises prices above marginal cost and allows suppliers to charge fees for moving data off platform, reducing the benefits from adoption.

I estimate the model using a newly-merged dataset on French firms during the period 2017-2021. I merge Orbis data on firms revenue and costs to a survey on firms' technology usage from Aberdeen, a market research firm. I manually construct a dataset of cloud computing prices and features from a set of first- and third-party websites. My data notably contains data on firms' cloud computing usage by supplier, allowing me to analyze bundling and switching over time. In addition, I observe features of firms' on-premise software stack and distance from cloud datacenters, such that I can construct a rich model of supplier choice.

Using my dataset, I first document overall patterns of usage. I find that adoption has risen rapidly, from around 12% of firms in 2017 to 28% by 2021. Bundling is common, with over half of cloud-using firms using multiple suppliers in 2021. Larger firms, and firms in more high-technology industries are more likely to adopt cloud computing. Despite rapid adoption, the upstream market remains concentrated - among firms in my sample, over 75% use two suppliers - Microsoft and Amazon.

I estimate my model in two stages. First, I estimate an industry-level model of production, using a control function approach to control for endogenous selection into cloud computing (de Loecker (2013)). I find that adopting cloud computing increases firm productivity by 0.2-1.5%, with the largest gains for high-technology manufacturing firms. I find support for the primary externality: bundling amplifies the productivity effects of cloud computing adoption. Over time, firms that use cloud computing have noticeably higher productivity levels than those that do not, and those that bundle suppliers experience additional productivity growth.

Second, I estimate the adoption, switching, and maintenance costs of using cloud computing using the full-solution approach to estimate a Markov moment equilibrium (Ifrach

and Weintraub (2017)). I show that adoption costs are large but falling over time due to global progress in reducing frictions. While switching costs are large, around €3 M per petabyte of data stored, I find that persistence in provider choice is driven by price as well as complementarities between cloud computing and on-premise software.

Using these estimates, I simulate three counterfactual scenarios, based on existing policies to promote cloud computing adoption. To simulate the impacts of market power on pricing exactly, I consider two scenarios - the addition of a new entrant and a horizontal merger between Microsoft Azure and Google Cloud Platform. I find that market concentration is harmful for total welfare, but at a declining rate - the addition of a new entrant increases total surplus by 2.82% annually and causes prices to fall by 30.7%, while the merger reduces welfare by only 0.29% per year (€905.3 million), even though prices rise by 25.9%. Next, I compare the impact of the two competition policies. I simulate an egress fee ban, or ban on charging consumers to move data off of a cloud provider's servers, based on the one implemented in the EU Data Act of 2024, which reduces firms' switching frictions based on their data stock and can lead to increased rates of adoption and bundling. I find that, holding supplier prices fixed, an egress fee ban leads to consumer surplus gains of 0.11% or €313.8 per year, while reducing supplier profits by 0.91% or €121.2 million per year. Second, I simulate increased complementarity between Microsoft server products and rival cloud platforms, based on concerns about anticompetitive policies raised by the UK Competition and Markets Authority. Such a mandate would increase welfare by 0.16% per year, or €517.2 million.

My results have implications for the design of competition policy for large technology platforms. First, it is important for regulators to consider the long-run effects of policies in expanding industries. For example, I find that while the short run an egress fee ban lowers consumer surplus, over the long run the policy increases consumer surplus by reducing the ability of suppliers in a mature market to lock in consumers. Second, policies should account for the impacts of consumers who multihome. As cloud platforms are not pure substitutes, policies leading to increased price competition do not necessarily reduce total surplus. For example, I find that the addition of a new entrant increases Microsoft's profits, as many Azure users adopt the new platform as a secondary cloud.

The role of switching frictions, bundling, and multihoming has informed competition policy in other ICT industries, including telecommunications (Park (2011); Viard (2007)) and web browsing (U.S. Department of Justice (2001)). The primary features of my framework–endogenous adoption and growth when consumers are firms, the evolution of an industry still in an expansionary phase, and dynamic supplier price responses–can be applied to study these other settings. In particular, as the rapidly expanding generative AI continues

to gather steam, the insights from my model may be used to design competition policies for this emerging general-purpose technology (cf. Korinek and Vipra (2025)).

Related literature: My paper contributes to several bodies of literature. First, it contributes to the literature on the dynamic productivity effects of firm technology adoption (Rubens (2025); Hornbeck et al. (2024); Humlum (2022); Aw et al. (2011)). I model cloud computing products as having both features of substitutes, as did water and steam power in Hornbeck et al. (2024), and of complements when firms multihome, as in Peters et al. (2017), who find complementarities between product and process innovation. Unlike prior literature, my model focuses on the impact of product differentiation and pricing on the downstream benefit of adoption.[1] I exploit the ability to observe adoption at the supplier level to model firms' costs and benefits of adoption as depending on prices charged by providers and a rich set of product features.[2]

Second, this paper contributes to the literature on the economics of IT adoption, particularly cloud computing, on firm productivity (De Ridder (2024); Brynjolfsson and Hitt (2003); Tambe and Hitt (2012); Bloom et al. (2016); Thompson (2017)), growth (Azar (2021); Barefoot et al. (2018); Bresnahan and Trajtenberg (1995)), and innovation (Besiroglu et al. (2024)). While prior work has focused on the productivity gains from a broad suite of cloud computing technologies (Jin and Bai (2022); DeStefano et al. (2023)), I focus particularly on infrastructure-as-a-service, the leasing of computation and storage.[3] My paper develops a new structural model of cloud adoption, featuring entry, switching, and multihoming decisions as well as provides the first, to my knowledge, empirical estimation of cloud supplier pricing conduct.[4]

Third, the present work contributes to the literature on the impact of market power in input markets on downstream firm outcomes. There is a large literature that has analyzed the role of input market power on downstream firm outcomes in macroeconomics (Acemoglu and Azar (2020); Baqaee and Farhi (2018)) and industrial organization (Yang (2020); Atalay et al. (2014); Crawford and Yurukoglu (2012)). Few other papers combine endogenous supplier pricing as well as dynamic capital accumulation and input choices by downstream firms. The paper closes to mine is Bertolotti et al. (2025), which analyzes a model of oligopolistic competition in the semiconductor industry, in which supplier market power serves as a

[1] Gortmaker (2024) focuses on how heterogeneity in open source web software impacts firms' production costs, but does not consider productivity benefits.

[2] More broadly, this paper contributes to a literature of dynamic discrete choice with switching frictions (for examples, see Honka (2014); Shcherbakov (2016); Polyakova (2016); Yeo and Miller (2018); Gowrisankaran and Rysman (2012); Dubé et al. (2009)).

[3] More broadly, this paper is also related to a wider literature on capital adjustment frictions (Asker et al. (2014); Bloom (2009)) and how leasing capital can reduce such frictions (Gavazza (2011)).

[4] Kilcioglu and Rao (2016) build a theoretical model of price and quality competition.

form of adjustment friction for downstream firms. Unlike this paper, mine features richer consumer dynamics such as switching costs and multihoming. In addition, I analyze the role of competition policies to promote consumer surplus.

Finally, this paper contributes to the literature on the implementation of competition policy in oligopolistic markets for general purpose technologies (Kearns (2024); Igami and Uetake (2020); Goettler and Gordon (2011)).[5] While the majority of these papers focus on policies that increase innovation, I model innovation as exogenous, since the measurable parts of quality improvements are a result of innovations further upstream by semiconductor (e.g. Intel) and hard-disk manufacturers (e.g. Western Digital). Instead, I focus on the implementation of policies that affect pricing. [6] A few theoretical papers focus on the impact of an data egress fee ban on cloud computing markets, finding that such a ban is either welfare reducing (Gans et al. (2023)) or depends on the demand elasticity (Do and Miklós-Thal (2025)). I estimate such a counterfactual and find that a ban is welfare-improving, though suppliers see a welfare decrease.

The rest of the paper proceeds as follows. Section 2 provides background on the cloud industry, focusing on the economic forces driving concentration in cloud computing and regulatory responses. Section 3 describes my dataset, including summary statistics and stylized facts that will drive my modeling decisions. Section 4 presents a structural model of firm production with cloud supplier choice, featuring multihoming and adoption and switching frictions. Sections 5 and 6 describes the my estimation procedure, including both firm production and cloud computing costs, sources of identification, and results, comparing them when possible to existing estimates in the literature or in industry reports. Section 7 considers my three counterfactual scenarios: the first two scenarios consider cases both holding supplier prices fixed and allowing suppliers to respond to the implemented policy. Finally, Section 8 concludes.

2 Background

Given that cloud computing is a technology with wide applicability, the welfare loss to consumers from upstream market power may be large. This section provides background on cloud computing, including adoption and technical details. I then discuss upstream market structure and regulatory responses to increase competition and consumer adoption. Further details are in Appendix F.

[5] More geenrally, this paper also contributes to the literature on the diffusion of GPTs (Humlum (2022); Chari and Hopenhayn (1991).

[6] Zhuo and Tiew (2025) focuses on the impact of location subsidies on the pricing of semiconductors.

2.1 Background on cloud computing adoption

Cloud computing allows firms to lease computation, storage, and software from cloud providers instead of investing in their on-premise IT services (Asker et al. (2014); Gavazza (2011); DeStefano et al. (2023)). Cloud firms invest in physical computing hardware and lease computing capital, as well as complementary software and services, to firms. This allows cloud providers to take advantage of economies of scale, spreading datacenter costs over a larger set of computing hardware in order to lower prices below that of on-premise computing.[7]

As a result of these productivity gains, cloud adoption is widespread. The European Commission reports that 45.2% of European firms have adopted the technology by 2023 (Eurostat / European Commission (2023)). Likewise, using Census data, Zolas et al. (2021) find that 78.4% of U.S. firms that use IT capital purchased at least one cloud service in 2018. Existing studies have found significant productivity gains from cloud services (DeStefano et al. (2023); Jin and McElheran (2024)).

Cloud computing comes in three primary forms: software-as-a-service (SaaS), platform-as-a-service (PaaS), infrastructure- as-a-service (IaaS). The first is the broadest and encompasses software leasing, such as Adobe Creative Cloud and Microsoft Office 365, hosted in the cloud. PaaS, such as Dropbox or Google Drive, also abstracts away from the underlying hardware by focusing on an application layer such as serverless computing or containers. IaaS is the simplest service, focusing on the leasing of "pure" computation (usually in the form of virtual machines (VMs)) and storage (usually in the form of "blobs"). For more detail, see Appendix F.1.

Since my data allows me to distinguish between IaaS, PaaS, and SaaS adoption, throughout the remainder of this paper, I focus on IaaS for two reasons. First, IaaS is the most popular cloud computing service.[8] Second, pricing is relatively straightforward, with consumers choosing from a menu of compute and storage products and paying a linear price per hour in the case of VMs or gigabyte-month in the case of storage. In addition, a significant majority of consumers are price-takers.[9] This allows me to use supplier-year average list

[7]Studies commissioned by cloud providers tend to find large economic impacts. A 2019 study commissioned by Microsoft finds a return on investment of 428% over 3 years North et al. (2019), with fixed migration costs totaling $1.38 million. A 2022 study of AWS finds an ROI of 241% over 3 years, with fixed costs totaling $604,801. Phillips and Owens (2022). Finally, a case study of an IT services firm in the UK oil and gas industry finds a 37% cost reduction from using Amazon's IaaS relative to on-premise and a 21% reduction in support calls from customers Khajeh-Hosseini et al. (2010). However, Emeras et al. (2016) find that usage of the University of Luxembourg's in-house computing cluster was less costly than leasing the equivalent compute on Amazon's AWS.

[8]See Brand et al. (2024) Appendix A. According to the UK Competition and Markets Authority, in 2022, Azure's top product was compute, or IaaS virtual machines (Competition and Markets Authority (2025a), 17).

[9]In a survey of IaaS users, the UK CMA found that only the largest 5% of firms were able to negotiate a

prices since my data do not contain consumer level prices.[10]

2.2 Market structure and regulatory response

Like many information technology markets with large returns to scale, the market for cloud computing services is highly concentrated, potentially allowing suppliers to raise prices above marginal cost and reduce consumer welfare. Worldwide, it is largely dominated by Microsoft, Amazon, and Google. The UK Competition and Markets Authority, for example, found that the HHI of the IaaS market was over 3000 between 2020 and 2023 (Competition and Markets Authority (2025h), p.90). In 2025, the top 3 cloud providers had a 70% market share in the IaaS market (Statista). This is due partly to the high fixed costs of entry (cf. De Ridder (2024)), which requires large datacenter, storage device, and semiconductor investments.[11]

As a result, cloud suppliers are able to price discriminate against consumers in the form of egress fees, or fees to move data out of a cloud service.[12] Figure 1 displays an example of these egress fees for Microsoft Azure. While these transfers have nonzero marginal cost, evidence suggests that firms charge a markup for data transfer (Do and Miklós-Thal (2025); Competition and Markets Authority (2025e), p. 9). In addition to shifting welfare from consumer surplus to producer surplus, such fees can also reduce the positive externality firms derive from multihoming, which I show below is an important driver of productivity gains from adopting IaaS.[13]

To improve competition, regulators have honed in on a few remedies - capping data egress fees and mandating software interoperability between IaaS suppliers. Since 2022, regulators in the UK, Japan, South Korea, France, and the European Union have considered banning data egress fees or limiting them to marginal cost (Korea Fair Trade Commission (2022)). In France, the SREN law (Loi Visant à Sécuriser et à Réguler l'Espace Numérique) entered into force on May 21, 2024. The law goes hand-in-hand with the EU Data Act, which will enter into force on January 12, 2027. Both regulations ban egress fees for cloud computing customers. On the other hand, the UK CMA recommends against capping egress fees at

price below list price (Competition and Markets Authority (2025a), p. 8). For further discussion, see Brand et al. (2024).

[10] For the remainder of the paper, I use the terms "IaaS" and "cloud computing" interchangeably.

[11] Google Cloud, for example, did not experience positive operating income until Q1 2023.

[12] Besides the egress charged by firms, firms also incur switching frictions from learning a new platform, including retraining employees, hiring new specialists or consultants for their provider(s) of choice, and rewriting software to use the new platform.

[13] The evidence for whether egress fees hinder multihoming is mixed. A 2022 study by the Korean Fair Trade Commission found that 60.3% of firms cited technological frictions and costs of data migration as a barrier to multihoming (Korea Fair Trade Commission (2022)). On the other hand, the UK CMA found that several mentioned such costs as a minor barrier to using multiple providers (Competition and Markets Authority (2025d) pp. 4, 7-10).

Internet Egress (routed via Microsoft Premium Global Network)

Source Continent	First 100 GB/Month	Next 10 TB/Month	Next 40 TB/Month	Next 100 TB/Month	Next 350 TB/Month
From North America, Europe to any destination	Free	$0.087 per GB	$0.083 per GB	$0.07 per GB	$0.05 per GB
From Asia (China excluded), Australia, MEA to any destination	Free	$0.12 per GB	$0.085 per GB	$0.082 per GB	$0.08 per GB
From South America to any destination	Free	$0.181 per GB	$0.175 per GB	$0.17 per GB	$0.16 per GB

Note: 1 TB = 1,000 GB

*Contact us for over 500 TB

Figure 1: Sample of data egress fees charged by Microsoft Azure in 2025.

marginal cost due to the technical difficulties of monitoring and enforcement and the fact that only a few firms they surveyed mentioned egress fees as a barrier to switching suppliers or exiting (Competition and Markets Authority (2025d) 4, 7-10). In response to the Data Act, all three of the major providers have eliminated data egress costs for exiting customers, but have not changed costs for continuing customers (Competition and Markets Authority (2025g), 44).[14]

A second remedy relates to Microsoft's vertical integration of popular server software–Windows Server and SQL Server–that makes it feasible to foreclose their use on rival cloud platforms. In particular, Microsoft prevents customers from transferring on-premise server software licenses to public cloud services, though Microsoft server software consumers can obtain a discount when purchasing Azure virtual machines using those software through its Auzre Hybrid Benefit program, launched in 2016 (Competition and Markets Authority (2025h)). As a result, the CMA has found the median consumer pays more to use Microsoft server software on Google or Amazon's clouds than on Microsoft's (Competition and Markets Authority (2025f)). This has led for calls to mandate price parity for Microsoft server licenses across cloud platforms.[15]

[14]In Appendix A.10, I exploit Google's unexpected announcement on January 11, 2024 that it would drop egress fees for exiting customers worldwide to see whether a drop in these fees reduced demand for Google virtual machines in the short run. I estimate a difference-in-differences specification 54 days around the announcement of the effect of the announcement on virtual machine spot prices, finding no significant effects. The estimation window is limited by the fact that both Amazon and Microsoft responded to Google's announcement by also dropping prices in the EU in early March 2024.

[15]For further discussion of the regulatory environment in France and worldwide, see Appendix F.2.

3 Data

To study the determinants of cloud computing adoption and its impact on firm production, I construct a dataset of firms' inputs, output, and technology usage, as well as cloud supplier prices, features, and costs. To obtain data on both sides of the market, I merge several data sources to create a novel dataset of firms' cloud computing supplier investments, non-IT investments, and prices. In this section, I describe the procedure used to match firms' financial data to technology usage, including data on cloud usage and other computing inputs. Then, I describe how I compute supplier prices, features, and marginal costs. Finally, I describe several stylized facts about cloud adoption and usage that will inform my empirical model. For more details on the data, see Appendix A.

3.1 Dataset construction

To build my dataset, I combine data from Orbis, the Aberdeen CI Technology Database (CITDB), anonymized IaaS usage data from Microsoft Azure, list prices from the three largest cloud providers and CPU quality and cost measures from Passmark. The panel covers firms at the annual level from 2015 to 2021.

Firm level financial data. I use data from the Orbis dataset for France from 2015-2021, including revenue, capital stock, materials expenditures, labor, and the wage bill. Due to French financial reporting requirements, Orbis covers a large fraction of total output of both public and private firms, limiting sampling biases. As such, data on France specifically has been widely used in the literature on firm-level productivity (e.g. Asker et al. (2014); Bloom et al. (2016)). While I only observe data on computing inputs for a subset of firms, I use the full Orbis sample to construct industry-location-year average wages and industry-year average depreciation rates at the 4-digit NAICS code level.

Firm technology adoption. I use the Aberdeen CITDB to obtain firms' cloud vendor choices, data storage and other IT hardware investments, and vendor choice for Microsoft server technologies. Aberdeen is a market research firm that seeks to obtain data on firms' IT usage, hardware, and software investments in order to sell to IT vendors, which it obtains largely through interviews, supplemented with interpolation and web-scraping. It has been used extensively on research studying the impacts of IT inputs on firm performance (Brynjolfsson and Hitt (2003)), including cloud computing (Bloom and Pierri (2018)). For a full description of the methodology, see Appendix A.1. The data is reported at the establishment level. To obtain firm-level data, I sum investment data across domestic establishments. Cloud computing usage and server vendor are reported at the extensive margin. I assume a firm uses a particular cloud provider if any of its establishments is observed having used the

provider.[16]

Cloud pricing. Cloud pricing data come from web scraping the archived VM pricing pages on Wayback Machine for Microsoft and Google. Amazon's data is not archived directly, so I use archives of EC2Instances.info Contributors (2025) to construct my panel. I scrape pricing for virtual machines at all datacenters for a given provider nearest to an adopting firm's headquarters at a quarterly level and then average across quarters. Since I do not observe which datacenters firms are using, I follow the assumptions in Wang et al. (2019) and assume that firms use the datacenter closest to their headquarters at the time of entry.[17] Figure A.3 in Appendix A.6.4 provides maps of suppliers' datacenter locations in 2017 and 2021. I also scrape data egress fees from Microsoft and Google's sites. As Amazon's data is not archived between 2015 and 2025, I use current prices and price change announcements on the AWS blog to reconstruct the price history (Amazon Web Services (2025)).

As there are a large number of SKUs across virtual machines and storage, I make two simplifying assumptions about the prices firms face. First, I assume that firms respond to VM prices over storage prices. Prior research has shown that firms that use IaaS use more computation than storage (Demirer et al. (2024)). In addition, Amazon's storage prices are not fully archived, making it difficult to reconstruct an accurate price history. Second, to deal with the high number of SKUs offered by each supplier, I take the weighted average of the price per vCPU hour at each datacenter-quarter-year, where the weights come from an anonymized representative sample of 5,000 customers of the monthly virtual machine usage at a Microsoft Azure datacenter in 2017 and again in 2019 (Cortez et al. (2017)).[18] For Microsoft Azure, I match prices to usage exactly. For Google and Amazon, I match each product to the nearest product in the Azure trace data, using fuzzy matching on product features such as core count and memory. Summary statistics, discussion of the match procedure, and further details are presented in Appendices A.2-A.7.

Other data sources. I use a number of auxiliary datasets to add control variables and deflate quantities. The data on datacenter locations for each cloud provider comes primarily from `datacenters.com` (Datacenters.com (2025)). For the remaining datacenters, I manually compile data from the colocation partners of the major cloud providers, and then obtain the latitude and longitude from `datacenters.com`. Second, I data from PassMark Software (2026) to compute CPU quality, which I take to be the weighted average benchmark

[16] See Appendix A.3 for a table of the IaaS technologies I observe for each provider.

[17] Some firms may choose datacenters that are instead near a majority of their consumers. See Wang et al. (2019) for a discussion of the attenuation bias that this produces. Demirer et al. (2024) provides evidence using internal data from a major cloud provider that datacenter choices are persistent, even within a supplier, and that 98% of firms headquartered in the US or EU use only datacenters in that region.

[18] Much of the variation in pricing across SKUs comes from the number of vCPUs offered on the processor. Dividing by vCPUs reduces the pricing variance.

score of CPUs in a datacenter, where weights again come from Cortez et al. (2017). Passmark also provides data on CPU prices to calibrate marginal costs. Third, data used to deflate the prices of inputs and outputs comes from Eurostat, with the exception of the price index for computer hard drive storage, which comes from the US Bureau of Labor Statistics (BLS).

Matching and cleaning. As the downstream firm datasets do not contain a common firm identifier, I merge them using a combination of exact and string-distance metrics based on firm name and other identifying information, such as website URL and zipcode. The match procedure is described in Appendix A.2. Using this procedure, I am able to match 24.2% of firms in my Orbis sample. To clean the data, I replace missing revenue and employment observations in the Orbis data with data from Aberdeen, which in turn sources these data from Dun & Bradstreet, and missing depreciation rates and wages with their averages at the industry (NACE 4-digit)-region (NUTS3)-year level. I deflate nominal values for computational investments, capital, materials, and output using industry-level Eurostat price deflators. I deflate storage investments using data from the US Bureau of Labor Statistics (McCallum and of Labor Statistics (2024)). Finally, I use the deflated book value of capital as well as Orbis depreciation data to construct a measure investment. If depreciation data is omitted, I use the average annual depreciation rate constructed from the entire Orbis sample.[19]

Summary statistics for the matched sample are displayed in Table A.2. Prior research has demonstrated that the response rate to Aberdeen is increasing in firm size (Bloom et al. (2016)). In Appendix A.4 I compare the full Orbis sample to my matched subsample, showing that the average firm in the matched sample is 41% larger than the average firm in Orbis, as measured by employment.

3.2 Cloud adoption patterns

To motivate my empirical model below, I present a series of stylized facts regarding cloud computing usage in France. The first two motivate my modeling choices. In particular, I demonstrate that the market is rapidly expanding, with a near doubling in usage over my sample period. Second, I show that the upstream market is highly concentrated, dominated by 3 large suppliers. The second two facts motivate recently-proposed pro-competitive remedies. First, I show that egress fees are negatively associated with firms' decision to switch suppliers. Second, I show that firms that use Microsoft server products have a greater probability of using Microsoft cloud products relative to those of other suppliers, indicating

[19] I use book values of capital rather than the perpetual inventory method due to data limitations, as Orbis only includes a full sample dating 10 years prior to the download date (January 2024). Details on the constriction of capital stocks and investment flows are in Appendix A.5.

Table 1: Summary Statistics

Variable	Observations	Mean	25th Percentile	Median	75th Percentile
Revenue	771,217	23.51	1.04	2.38	7.09
Labor	771,217	70.83	6.00	14.00	33.00
Capital	762,623	28.28	0.12	0.41	1.37
Materials	521,539	22.88	0.78	2.28	8.51
Computation	771,217	77.12	2.00	10.00	28.00
Firm Age	771,217	23.06	11.00	20.00	31.00
Cloud Usage	771,217	0.11	0.00	0.00	0.00

Notes: The units for revenue, capital, and materials are in millions of real 2015 euros. The units for labor and computation are in units. The unit for firm age is years.

product complementarities.

Fact 1: Cloud computing usage is rising rapidly, and purchasing from multiple suppliers is common. In Figure 2, I plot the fraction of firms in my sample using cloud computing across all years.[20] Growth has been rapid, increasing from 12.4% of firms in 2017 to 27.9% in 2021, with a rapid increase during the COVID-19 pandemic. In addition, a large minority of firms are using two or even three cloud providers. On average, 45.95% of cloud users in my sample use more than one cloud provider, and 23.28% use all three major cloud providers.[21] As a result, several non-leading cloud providers, including Google and IBM, have sought to position themselves as the optimal second cloud for consumers (Competition and Markets Authority (2025a), 23).[22]

Industry case studies show that firms experience strong complementarities from using multiple cloud providers. Multihoming allows firms to choose suppliers on whose virtual machines software or tasks run most efficiently. For example, in 2018, Capital One began using both Amazon and Google clouds - leveraging Amazon for its reliable infrastructure and Google for its data analytics software (MultiCloud.Tips (nd)). In addition, firms may use multiple providers to enhance reliability in the case of outages, as in the case of Adobe. However, multihoming has downsides. Besides the switching costs mentioned above, managing

[20] Aberdeen only tracks IaaS usage in 2016 for establishments that survived until 2020. To avoid this survivorship bias, I begin the figure in 2017.

[21] In the UK, the rates of multihoming among cloud users was 7.04% in 2020 and 7.50% in 2021 (Competition and Markets Authority (2025b), 5). On the other hand, Flexera finds that 86% of surveyed firms use more than one public cloud provider (Flexera (2020), 16.

[22] The overall level of multiclound integration within multihoming firms remains an open question. Cloud security firm Wiz found that in 2023, 2/3 of multihoming customers had over 80% of their data on the servers of a single provider (Piper and Cohen (2023)). Meanwhile, Seagate Technology (2020) finds that 48% of multicloud firms have applications across providers that interact, and 26% of multicloud firms have ported an application from one IaaS provider to another. Finally, Competition and Markets Authority (2024) considers it representative for a multicloud firm to backup all incoming consumer data on a secondary platform.

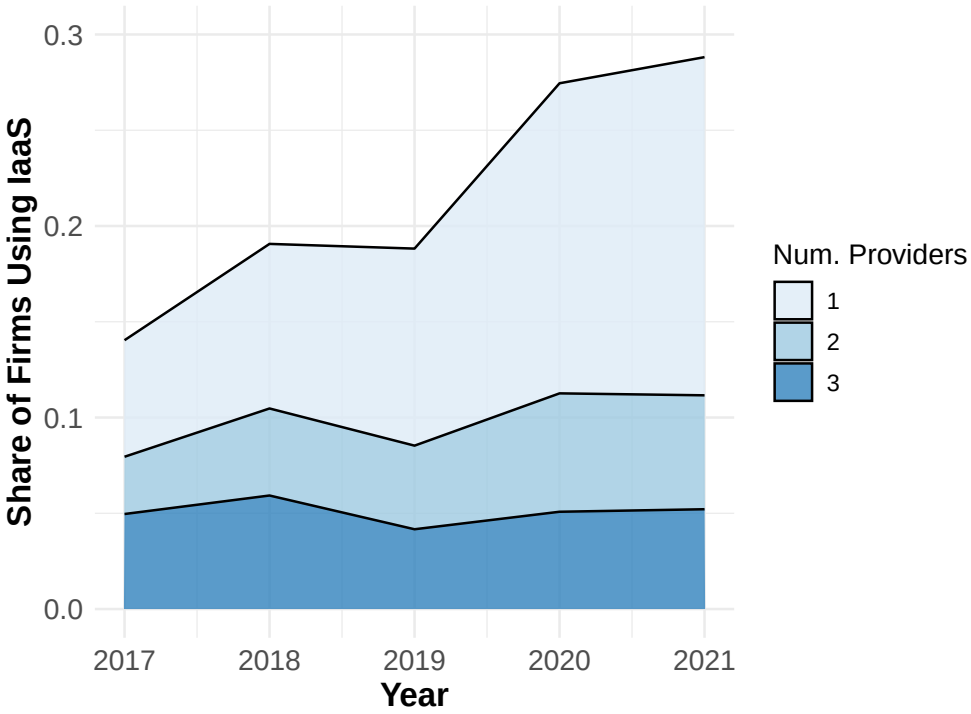


Figure 2: Growth of cloud computing usage over time

a workflow across multiple providers also requires additional costs - fixed costs of ensuring software compatibility and human capital to manage a workflow across multiple providers. For both of these reasons, one should expect to see larger firms be more likely to adopt multiple providers. Table 2 shows that multihoming firms tend to be larger than single-cloud users, with a higher intensity of IT usage, as measured by software investment.

Table 2: Comparison of single- and multihoming firms

	Med. Revenue	Med. Computers	Med. Employees	Med. Software Intensity
Single supplier	6.06	29	32	0.00448
Multiple suppliers	50.7	294	294	0.00498

Notes: Values computed among cloud adopters. Revenue is in million nominal euros. Software intensity is nominal software investment over revenue.

Fact 2: Supply of cloud services is highly concentrated The market for cloud services is dominated by three major firms - Amazon, Microsoft, and Google, with smaller providers such as IBM and Alibaba targeting specific industries. As of 2025, the three largest firms control 70% of the IaaS market, with Alibaba and IBM controlling another 18% (Statista (2025)). Figure 3 displays the market shares of the three largest suppliers in my data,

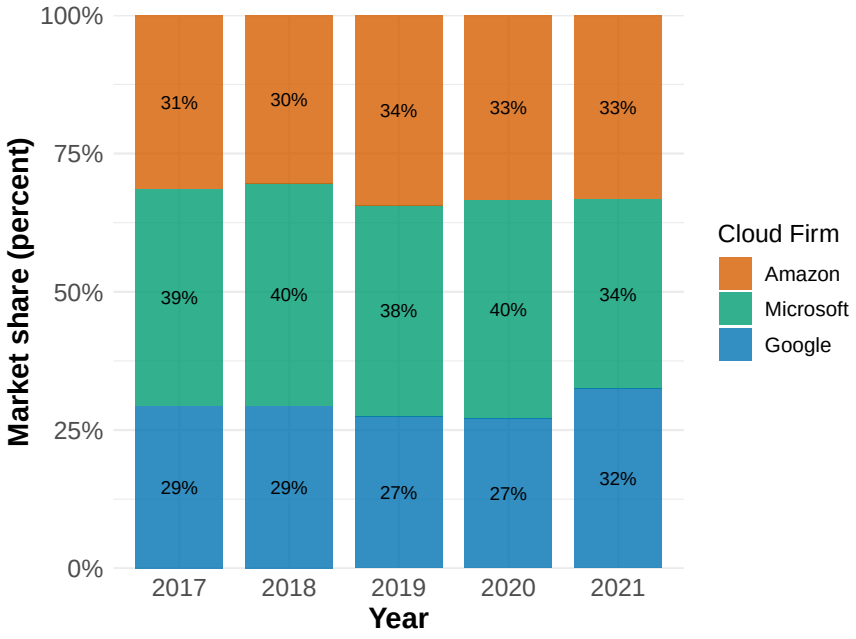


Figure 3: Cloud provider market share

Shares are weighted by firm software investments. For multihoming firms, consumption is divided evenly among suppliers used.

weighted by firm software investments.[23] Market shares are roughly constant over the sample period, with Microsoft slightly ahead of Amazon.[24]

The reasons for this dominance are numerous. As the first firm to offer cloud computing in 2006, Amazon has enjoyed a first-mover advantage that has persisted due to continued quality improvements as well as the lock-in effects highlighted above. As a result, the suppliers have enjoyed significant profits. Figure 4 displays the quarterly operating income, in real USD, of each supplier's cloud segment, starting in Q1 2013 for Amazon and Microsoft and Q4 2020 for Google.[25] While Google posted operating losses for the first few quarters for which data is available, by 2023, all three firms enjoy significant margins.

Fact 3: Cloud supplier choice is persistent During my period of study, innovation in the cloud computing industry has been rapid. All three firms significantly increased the number of regional datacenters (see A.3 in the appendix). Likewise, the number of virtual machine offerings increased rapidly between 2017 and 2021, including the addition of specialized processors, such as Google's tensor processing units (TPUs) for machine learning, which launched in 2018. Despite this innovation, the of switching cloud supplier bundles was low,

[23] Aberdeen CITDB also contains data on IBM IaaS usage for a small number of firms, but none are matched to my Orbis sample. Appendix A.9 presents resutls for the worldwide IaaS market share in 2023, including smaller suppliers.

[24] By 2025 Microsoft had 27% market share in the IaaS market compared to 20% for Amazon (Statista (2025).

[25] Note that this segment incorporates both IaaS and PaaS products.

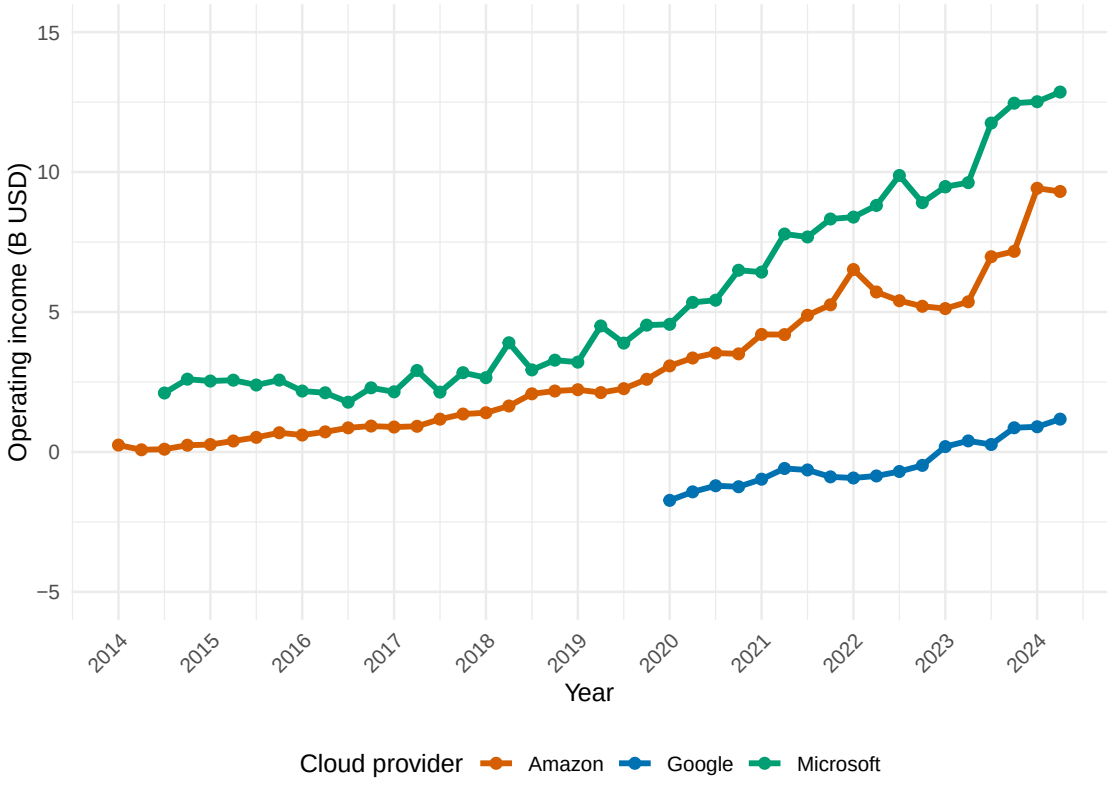


Figure 4: Operating income of the cloud divisions of the 3 largest cloud suppliers (Q1 2014-Q1 2024)

Table 3: Correlates of supplier switching

Variable	Model 1	Model 2
Intercept (ζ_0)	11.498 (0.452)***	11.510 (0.453)***
Log egress fee (ζ_1)	-0.816 (0.024)***	-0.816 (0.024)***
Log capital (ζ_2)		0.003 (0.022)
Log capital × log egress fee (ζ_3)		0.000 (0.005)
Log VM price (ζ_4)	0.136 (0.044)**	0.136 (0.044)***
Log VM quality (ζ_5)	-9.448 (0.297)***	-9.455 (0.297)***
Log Azure × Microsoft server usage (ζ_6)	0.021 (0.018)	0.020 (0.018)
Observations	80,830	80,830
Year FE	✓	✓
Industry FE	✓	✓
Log likelihood	-20321.86	-20316.56

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions estimated via a probit model with year and industry fixed effects.

averaging 2.64% of cloud users in my sample.[26]

To study the determinants of switching provider bundles, I run the following probit regression on the subset of firms that could incur an egress fee:[27]

$$\mathbb{P}[b_{ft} \neq b_{ft-1}] = \zeta_0 + \zeta_1 e_{ft}^{TB} + \zeta_2 k_{ft} + \zeta_3 e_{ft}^{TB} \cdot k_{ft} + \mathbf{X}'_{ft}\boldsymbol{\zeta}_{4:6} + \tau_i + \tau_t + \epsilon_{ft}$$

where the outcome is the probability that a firm f switches from a cloud computing bundle b_{ft}, where $b \in \{\text{Amazon}, \text{Microsoft}, ..., \{\text{Amazon}, \text{Microsoft}\}, ...\}$, e_{ft}^{TB} is the log of the cost faced by a firm to migrate 100 terabytes of data in millions of real euros. To control for the fact that fees scale with firm data storage, I use k_{ft}, the log of the firm's real capital stock. To control for observed factors affecting switching, I include $\mathbf{X}'_{ft}$ as a vector of controls - the average price per core-hour of b_{ft-1} in real euros, the average log VM quality, and an indicator for whether a firm used Azure and Microsoft Server products in period $t-1$. I also include industry and year fixed effects. Results are displayed in Table 3.

In both specifications, I compute an average marginal effect of -0.11, indicating that a 1% increase in egress fees reduces the probability of switching by 0.11%. Likewise, the coefficients on price and quality are significant and of the correct sign. This suggests that, controlling for observable product characteristics, higher switching costs are more likely to deter cloud users from switching suppliers.

[26] A 2023 Ofcom report found that approximately 20% of firms have ever switched cloud providers (Ofcom (2023).

[27] That is, firms using cloud computing in period t or $t-1$.

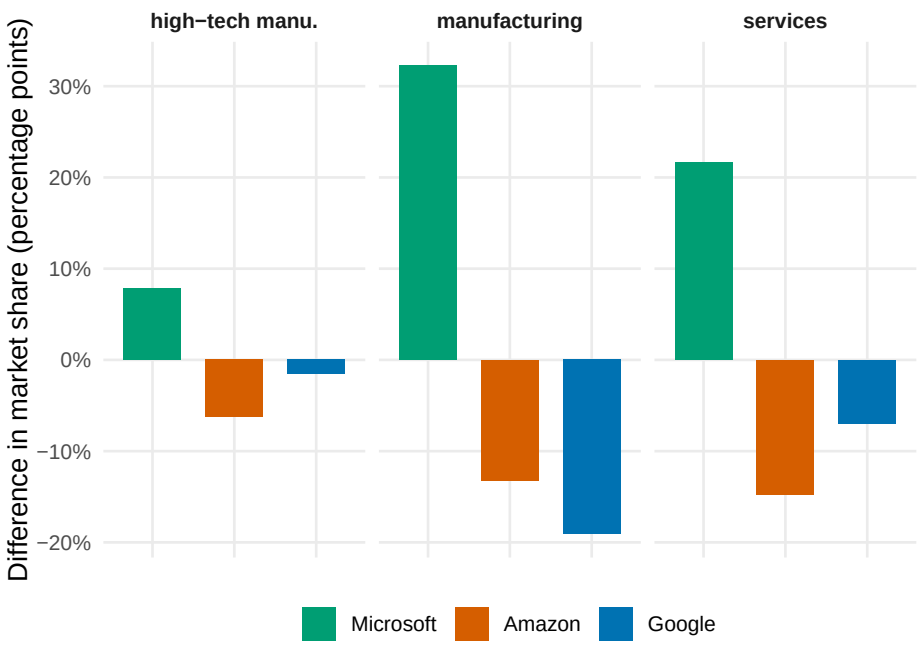


Figure 5: Difference in cloud supplier market share by Microsoft server adoption status

Fact 4: Server supplier predicts cloud supplier choice.

The dominant cloud suppliers are all large, multi-product firms that offer complementary goods to their cloud services. Arguably the most important of these is Microsoft, which is the dominant firm in the server software market and offers lower prices as well as other interoperability measures between its server and cloud products, as discussed above. In addition, Microsoft sales representatives often attempt to sell Azure to existing customers.[28] Figure 5 displays firm-level cloud supplier choice shares for firms with and without Microsoft Server products for firms in high-technology manufacturing (sectors with mean R&D expenditures above 5%), low-technology manufacturing, and services. Across sectors, firms that are using Microsoft Server products are much more likely to use Microsoft Azure as well.

3.3 Implications of empirical facts

The above four facts inform my choices in modeling cost heterogeneity in cloud adoption. The first fact suggests that the fixed costs of adoption, relative to on-premise computing, are rapidly falling, prompting a model of firms' dynamic decisions as a nonstationary equilibrium. In addition, the frequency of multihoming suggests that firms derive significant benefits from complementarities between cloud suppliers' products. I rationalize this by allowing firms to

[28] I thank Bill Whyman and the director of cloud infrastructure at a health insurance firm for independently making this observation.

derive additional productivity benefits from multihoming while facing additional variable costs from the organizational complexities of integrating multiple suppliers. Facts 3 and 4 suggest that the potential barriers to competition– egress fees and complementarities between IT vendors' products–are associated with consumers' choice of cloud supplier, highlighting the potential for welfare gains from policies that increase price competition among suppliers, analyzed in Section 7.

4 Model

To estimate the impact of upstream concentration on firms' cloud adoption and usage decisions, I develop a model of firm production with computational input choices. Firms produce output using a Cobb-Douglas production technology, and choose capital investment and computing bundles each period to maximize the net present value of current and future profit flows. Cloud adoption increases the future path of productivity, and multihoming generates additional productivity benefits. In the short run, firms pay fixed costs of adoption or bundle switching. Firms are price-takers, choosing the suppliers with the best price and feature set to minimize variable costs of using cloud computing. To study supplier pricing incentives, I construct a model of Betrand competition in which suppliers observe the vector of market shares and choose prices and data egress fees to maximize profit.

4.1 Firm production and cloud computing demand

I first describe how cloud computing affects firms' productivity and output. Firms make a discrete choice of a bundle of computing inputs $b_{ft} \in \mathcal{B}$. Firms can choose to use on-premise computing only, $b_{ft} = \varnothing$, or on-premise computing and any number of cloud providers in my sample: $b_{ft} \in \{\text{Amazon}, \text{Google}, \text{Microsoft}, \{\text{Amazon}, \text{Google}\}, ...\}$. Including the on-premise-only bundle, $b_{ft} = \varnothing$, there are 8 discrete choices: $|\mathcal{B}| = 8$. Adopting cloud computing affects firm productivity, $\Omega_{ft}(b_{ft})$. In addition to a bundle of computing providers, firms choose static inputs of labor and materials, L_{ft} and M_{ft}, as well as a dynamic capital input, K_{ft}. In each industry $i \in \mathcal{I}$, firms produce output according to the following Cobb-Douglas production technology:[29]

$$Y_{ft} = \Omega_{ft}(b_{ft}(\mathbf{X}_{fbt}))L_{ft}^{\beta_l}M_{ft}^{\beta_m}K_{ft}^{\beta_k} \tag{1}$$

[29]I employ a Cobb-Douglas specification in order to simplify estimation of the dynamic model. As well, a Cobb-Douglas production function is the first-order approximation of any general production function. In Appendix D.1.2, I estimate a translog production function, finding similar results.

where Y_{ft} is revenue; β_X for $X \in \{L, M, K\}$ are the output elasticities; and Ω_{ft} depends on the chosen IT input bundle through its feature vector $\mathbf{X}_{fbt}$. Throughout the remainder of this paper, I will focus on the above revenue production function and its associated notion of productivity, TFPR.[30] Assuming that firms take input prices p_{Xt} as given, I can use the firm's cost minimization assumption to solve for optimal L, M and rewrite its flow profit net of cloud computing costs in terms of the state variables (Ω_{ft}, K_{ft}) as

$$\pi_{ft}(\Omega_{ft}, K_{ft}) = \Lambda \Omega_{ft}^{1/(1-\beta_L-\beta_M)} K_{ft}^{\beta_K/(1-\beta_L-\beta_M)} - C^K(I_{ft}, K_{ft}) - \kappa(b_{ft}) D_{ft} \tag{2}$$

$$\text{where } \Lambda = (1-\beta_L-\beta_M) \left(\frac{\beta_L}{p_{Lt}}\right)^{\beta_L/(1-\beta_L-\beta_M)} \left(\frac{\beta_M}{p_{Mt}}\right)^{\beta_M/(1-\beta_L-\beta_M)} \tag{3}$$

where I_{ft} is the period t investment in K_{ft}, and $C^K(\cdot)$ is a convex adjustment cost function parameterized as as follows:[31]

$$C^K(I_{ft}, K_{ft}) = p_t^K I_{ft} + C_1^K K_{ft} \left(\frac{I_{ft}}{K_{ft}}\right)^2 - I_{ft} \epsilon_{ft}^K \tag{4}$$

where $\epsilon_{ft}^K \sim \mathcal{N}(0, \sigma_K)$ is an i.i.d. shock to the cost of investment incurred by a firm, rationalizing within period variation in investment (cf. Gowrisankaran and Schmidt-Dengler (2024)). Next, $\kappa(b_{ft})$ refers to the total cost of cloud computing adoption and maintenance, which are scaled by a firm's data stock, D_{ft}. Finally, I assume K_{ft} evolves according to the law of motion

$$K_{ft} = (1-\delta_{it}) K_{ft-1} + I_{ft} \tag{5}$$

for industry-year depreciation rate δ_{it}.

Productivity process. Next, I let ω_{ft} follow the industry-specific law of motion given by the AR(1) process:

$$\omega_{ft} = g(\omega_{ft-1}, b_{ft}(\mathbf{X}_{bft}), \mathbb{1}\{t = 2020\}) + \nu_{ft} \tag{6}$$

where $\nu_{ft} \sim \mathcal{N}(0, \sigma_i^\omega)$ is an i.i.d. shock. I assume that a cloud computing bundle choice affects firm productivity according to the function $g(\cdot)$, which is dependent on past productivity, cloud bundle choice, and a COVID-19 shock. In my empirical specification, I

[30] See Foster et al. (2008) for a discussion of the difference between TFPQ and TFPR and the limitations of the latter.

[31] Some models Cooper and Haltiwanger (2006); Bloom (2009) allow for a fixed adjustment cost as well, though Asker et al. (2014) finds it is precisely 0 in the French manufacturing sector.

find that only the number of providers, $\#j$, is statistically significant across sectors, such that I can parametrize the productivity process as

$$\omega_{ft} = \gamma_0 + \gamma_1 \omega_{ft-1} + \gamma_2 \omega_{ft-1}^2 + \gamma_3 \omega_{ft-1}^3 \tag{7}$$

$$+ \gamma_4 \cdot \#j + \gamma_5 \mathbb{1}\{t = 2020\} + \nu_{ft} \tag{8}$$

This specification follows existing models of firm input adoption, such as Peters et al. (2017); Aw et al. (2011). While richer interactions between productivity and cloud usage are possible, I choose this simpler specification to keep the model tractable for my relatively short panel of data (cf. Humlum (2022)).

Computing bundle choice. Firms also face a dynamic decision to adopt a bundle of cloud computing providers b_{ft} or only adopt on-premise computing, $b_{ft} = \varnothing$. As with other capital goods, K_{ft}, firms make an adoption decision in period t, which is realized in period $t + 1$. Adopting and using cloud computing incurs both fixed and variable costs. Fixed costs are divided into adoption costs, κ_t^E, which are incurred in period t when a firm adopts cloud computing, and switching costs, κ^S, which are incurred when a firm switches cloud computing bundles or exits the cloud. The one-period time to build assumption and fixed costs represent the one-time costs of hiring, software rewriting, and data migration.[32] I allow the cost of entry to vary to account for the development of cloud migration tools and expansion of the cloud migration specialist labor force over time.[33] To reduce model complexity, I fix switching costs since I do not observe high variation in switching rates over time. For each period using a cloud bundle, firms incur a maintenance cost $\kappa_M(\mathbf{X}_{fbt})$ which varies with bundle characteristics $\mathbf{X}_{fbt}$. These costs represent the variable costs of deploying virtual machines and storage, in addition to the ongoing costs of employing software engineers with the relevant human capital and of ensuring software compatibility across multiple suppliers. For example, a firm using multiple providers is likely to have higher costs of ensuring continual compatibility of software across cloud services, as well as data egress fees for ongiong data transfers between providers. Finally, I assume that these costs scale with a firm's stock of data, a complementary good to computing (Brand et al. (2024)). [34] To

[32] While individual software developers may create VM instances for very little fixed costs, migration of existing firm workloads is costly. Since the Aberdeen data is sourced from interviews regarding firm production processes to sell to IT hardware and software providers, it is likely capturing the latter use case.

[33] Each of the cloud providers in my data offers an internal migration tool. In addition, firms can use third party migration software such as Cortex or Carbonite Migrate.

[34] While Aberdeen contains a number of variables related to firms' IT capital investments, I choose to use their data on storage investment rather than PC investment to proxy for investments in computation and data storage. Computer storage devices are relatively homogeneous (Igami and Uetake (2020)), while PCs

maintain tractability for estimation, I assume that data is part of a firm's total capital stock and, as such, there exists a one-to-one mapping between capital and data, such that I can write data as $D_{ft}(K_{ft})$.

Using these assumptions, I can write the cloud computing usage costs as follows:

$$\kappa_t = \underbrace{\kappa_t^E \cdot \mathbb{1}\{b_{ft+1} \neq b_{ft} | b_{ft} = \varnothing\}}_{\text{Entry cost}} + \underbrace{\kappa^S \cdot \mathbb{1}\{b_{ft+1} \neq b_{ft} | b_{ft} \neq \varnothing\}}_{\text{Switching friction}} + \underbrace{\kappa^M(\boldsymbol{\alpha}' \mathbf{X}_{bft}) \cdot \mathbb{1}\{b_{ft+1} \neq \varnothing\}}_{\text{Maintenance cost}} \tag{9}$$

where entry costs are parametrized as:

$$\kappa_{ft}^E = \exp(\kappa_0^E + \kappa_1^E \cdot t) + \epsilon_{ft}^E \tag{10}$$

where $\epsilon_{ft}^E \sim Gumbel(0, \sigma^E)$. The common component of the entry cost follows the standard functional form assumption in other models of firm adoption of technologies with rapidly falling costs (cf. Rubens (2025); Humlum (2022)). Denoting $server_{ft} \in \{0, 1\}$ as an indicator for whether firm f is using Microsoft server products in year t and P_{bft} as the mean price per core hour of a virtual machine in bundle b, maintenance costs are parametrized as

$$\kappa_{ft}^M = \alpha_0 + \alpha_1 \underbrace{P_{bft}}_{\substack{\text{Price} \\ \text{per core hour}}} + \alpha_2 \cdot \underbrace{\#j \cdot \mathbb{1}\{\#j > 1\}}_{\text{Multihoming}} + \alpha_3 \underbrace{\log(q_{bft})}_{\text{VM quality}} \tag{11}$$

$$+ \alpha_4 \underbrace{\mathbb{1}\{\text{Microsoft} \in b_{ft} \cap server_{ft} = 1\}}_{\substack{\text{Microsoft} \\ \text{complementarity}}} + \underbrace{\sum_{j \in b} \xi_j}_{\substack{\text{Vertical} \\ \text{differentiation}}} + \epsilon_{ft}^b \tag{12}$$

These maintenance cost shifters capture the baseline costs of using cloud computing relative to on-premise as well as how these costs vary with cloud computing prices, multihoming choices, virtual machine quality, complementarities with on-premise software stacks, and unobserved vertical differentiation.[35] Finally, $\epsilon_{ft}^b \sim Gumbel(0, \sigma^b)$ is an i.i.d. shock that firms in each period in which they are using cloud computing.[36]

are likely to differ on many other features in addition to raw computing power, which would introduce measurement error.

[35] I model Microsoft server adoption status as an exogenous AR(1) process to account for this complementarity, as described in Appendix C.1.

[36] It is well known that one cannot identify both maintenance costs and unobserved scrap values using only firm exit rates (Aguirregabiria and Suzuki (2014)). I choose to normalize scrap values to 0 because cloud provider features are more naturally modeled as maintenance cost shifters and not scrap value shifters.

Each period, then, firms observe their current productivity level and, conditional on their current capital and bundle choices, make capital investment and bundle decisions, which are realized in the next period. I assume that firms can perfectly forecast exogenous aggregate states - non-IT input prices, Microsoft server adoption decisions, and the growth rate in VM quality, which I assume is exogenously related to Moore's law.[37] For cloud computing prices, chosen by suppliers, I assume that firms forecast prices using an AR(1) decision rule, described below. Them, using $\mathbf{P}_t$ to represent the vector of input prices and quality, I can write the payoff relevant state as $s_{ft} = (\Omega_{ft}, K_{ft}, b_{ft}, server_{ft}, \mathbf{P}_t) \in \mathcal{S}$. I can then write firm value functions as follows:

$$V_{ft}(s_{ft}) = \max_{I_{ft}, b_{ft}} \pi_t(s_{ft}; \boldsymbol{\beta}, \boldsymbol{\gamma}) \tag{13}$$

$$- C^K(I_{ft}, K_{ft}) - D_{ft}\kappa^E_{ft} - D_{ft} \cdot \kappa^M_{ft}(s_{ft}; \boldsymbol{\alpha}) \tag{14}$$

$$+ \beta \mathbb{E}_t[V_{ft+1}(s_{ft+1})|s_{ft}] \tag{15}$$

where β is the discount factor. The time subscript t is due to the fact that profits π_{ft} will vary over time due to falling adoption costs and changes in input prices. In each period, each firm 1) draws relevant cost shocks for investment and cloud entry and maintenance costs, $\epsilon^K, \epsilon^E, \epsilon^b$, and decides its level of capital investment, I_{ft}, computing bundle to maximize the net present value of future profits.[38]

4.2 Supplier competition

Upstream, cloud computing suppliers $j \in \{\text{Amazon, Microsoft, Google}\}$ engage in Bertrand price-setting of virtual machine prices and data egress fees in the next period. Due to adjustment frictions, these prices are be a function of the current state, s_t. Before downstream firms make their investment choices in each period, upstream firms choose prices P_{jt} to maximize their value function:

[37] This is an upper bound on progress, since cloud firms choose to continue offer instance types based on older CPUs. I abstract away from quality choice in the model, however.

[38] The reasons for including a cost shock on maintenance costs are twofold. First, since I estimate the model on broad industries yet only observe investment cost at the annual level, it is likely there is meaningful empirical variation in the costs of investment for firms. Second, Gowrisankaran and Schmidt-Dengler (2024) demonstrate that the algorithm has speed improvements over transitional methods of finding firms' investment policy functions.

$$W_{jt} = \max_{P_{jt}} \mathbb{E}_t[\pi^S_{jt} + \beta W_{jt+1}(s_{t+1})|s_t] \tag{16}$$

$$\text{where } \pi^S_{jt} = (P_{jt} - C^{vm}_j)Q_{jt} + (E_{jt} - C^{egress})[Q^{exit}_{jt} + \lambda Q^{multi}_{jt}] \tag{17}$$

The price P_{jt} refers to the average price per vCPU-hour supplier j sets in a given year. While a richer model would consider the full menu of virtual machine prices, in practice, publicly-available virtual machine usage data suggests that variance in prices per vCPU-hour is rather small.[39] In addition, egress fees E_{jt} are set to deter consumer exit. In reality, I observe few nominal price changes in egress fees during my sample period, which are declining purely as a function of the inflation rate. Therefore, for computational tractability, I abstract away from egress fee choices by assuming suppliers always choose the observed egress fee schedule and then choose P_{jt} to maximize profit. Then Q_{jt} is the total quantity of data stored in the cloud,[40] while Q^{exit}_{jt} is the quantity of data that exits cloud provider j, and Q^{multi}_{jt} is the quantity of data using j's services among multihomers. I assume that, annually, multihoming consumers move a fraction $\lambda \in [0, 1]$ of their data between suppliers, incurring egress charges.[41]

It is natural to assume that suppliers are forward looking - in a market with rapid expansion, they are likely to price below their static marginal cost plus markup in order to grow their installed base. To preserve computational tractability, I represent suppliers' preferences for investing over harvesting using a semi-reduced-form conduct parameter, ζ_{jt} (cf. Bresnahan (1982)). That is, taking supplier j's first-order conditions gives

$$Q_{jt} + (P_{jt} - C^{vm}_j) \cdot \frac{\partial Q_{jt}}{\partial P_{jt}} + (E_{jt} - C^{egress}) \cdot \left[\lambda \frac{\partial Q^{multi}_{jt}}{\partial P_{jt}} + \frac{\partial Q^{exit}_{jt}}{\partial P_{jt}}\right] + \beta \frac{\partial \mathbb{E}_t W_{jt+1}(s_{t+1})}{\partial P_{jt}} = 0 \tag{18}$$

Rearranging gives

[39] See Appendix A.6 for a full analysis.

[40] This implicitly assumes that consumers use data and virtual machines in fixed proportion.

[41] Industry and regulatory reports find that multihoming firms tend to seek to minimize cross-cloud data transfer. In the baseline model, I set $\lambda = 0.15$, slightly lower than the conservative scenario analyzed in Competition and Markets Authority (2024), 132.

$$P_{jt} = \underbrace{C_j^{vm}}_{\text{marginal cost}} + \underbrace{\frac{Q_{jt}}{-\partial Q_{jt}/\partial P_{jt}}}_{\text{Markup}} + \underbrace{\frac{(E_{jt} - C^{egress}) \cdot \left[\lambda \frac{\partial Q_{jt}^{multi}}{\partial P_{jt}} + \frac{\partial Q_{jt}^{exit}}{\partial P_{jt}}\right]}{-\partial Q_{jt}/\partial P_{jt}}}_{\text{Egress revenue adjustment}} + \underbrace{\beta \frac{\partial \mathbb{E}_t[W_{jt+1}(s_{t+1})]/\partial P_{jt}}{-\partial Q_{jt}/\partial P_{jt}}}_{\text{Invest-harvest incentive}} \tag{19}$$

which can be rewritten as follows:

$$P_{jt} = C_j^{vm} + \zeta_{jt} \cdot \frac{\tilde{Q}_{jt}}{-\partial Q_{jt}/\partial P_{jt}} \tag{20}$$

$$\text{where } \tilde{Q}_{jt} = Q_{jt} + (E_{jt} - C^{egress}) \cdot \left[\lambda \frac{\partial Q_{jt}^{multi}}{\partial P_{jt}} + \frac{\partial Q_{jt}^{exit}}{\partial P_{jt}}\right] \tag{21}$$

$$\zeta_{jt} = 1 + \beta \frac{\partial \mathbb{E}_t[W_{jt+1}(s_{t+1})]/\partial P_{jt}}{\tilde{Q}_{jt}} \tag{22}$$

Thus, ζ_{jt} represents a measure of "invest vs. harvest" dynamics. If $\zeta = 1$, the supplier is focused on harvesting the existing user base, while if $\zeta < 1$, the supplier is setting prices below what static pricing would imply in order to expand its customer base. In the empirical application, this parameter will be estimated, allowing me to preserve computational tractability on the supply side via a simple solve of static first-order conditions each period while allowing for more realistic supplier pricing than static Bertrand would imply.

A few additional comments are in order. First, by focusing on price setting, I abstract away from dynamic investment decisions in improved products, such as new datacenter capacity (Pan Fang and Greenstein (2025)), or the development of custom processors, such as Google's tensor processing units (TPUs) or Amazon's Graviton series.[42] While a full treatment of these impacts is beyond the scope of this paper, I argue that Bertrand price setting accurately captures short run changes in virtual machine pricing. Prior research has found, for example, that Amazon and Microsoft responded rather quickly to Google's entry into the market by cutting prices (Byrne et al. (2018)).

Second, I note that ζ_{jt} is a parameter that collapses several forces–innovation, changes in demand, or competition from fringe firms–that can affect supplier pricing (cf. Corts (1999); Reiss and Wolak (2007)). Therefore it is not a primitive, but a reduced-form representation of supplier conduct.

[42] Vantage, a cloud computing FinOps platform, for example, found that Graviton-series VMs were less than 1% of spending among their AWS-using customers in 2022 (Vantage (2022)).

4.3 Timing and equilibrium

Finally, I describe the timing of the game and my equilibrium concept. With many firms selling heterogeneous products, for any individual firm, tracking the states of all competitors is infeasible. Given the number of downstream firms, solving for a MPE using Ericson and Pakes (1995) is computationally infeasible. Therefore, I adopt the Markov moment equilibrium (MME) concept from Ifrach and Weintraub (2017) as a simplifying assumption. In such an equilibrium, firms keep track of their own state, $\tilde{s}_{ft} = (\Omega_{ft}, K_{ft}, b_{ft}, server_{ft})$, as well as moments of the aggregate industry state. The payoff relevant states include the exogenous price changes in industry level wages, materials price, capital cost, and VM quality index $(P_i^l, P_i^m, P_i^k, \log(q_j))$ as well as the endogenous price of each supplier, P_t^b. I make the simplifying assumption that that firms forecast a simple AR(1) pricing rule (cf. Garcia-Osipenko et al. (2025); Hornbeck et al. (2024)), which I find works well empirically.[43]:

$$P_{jt+1} = \rho_0^j + \rho_1^j P_{jt} \tag{23}$$

Thus, downstream firms best respond to their own state as well as their perceived price forecast for each supplier. Finally, given the falling entry costs, this setting is nonstationary. Therefore, assume that the industry transitions from $t_0 = 2017$ to some period T, after which the industry is in a steady state. Thus, firms' strategies depend on the initial state, s_0. Throughout, I focus on symmetric pure strategies. The timing of the game is as follows:

1. The game begins in period $t = 2017$, at an initial industry state s_0.
2. Suppliers observe $\mathbf{s}_t$ and choose P_{jt}.
3. The idiosyncratic shocks σ^E, σ^B are realized and downstream firms choose b_{ft+1}
4. Investment cost shocks σ^K are realized and downstream firms choose investment I_{ft}.
5. Firms and suppliers receive per-period profits net of any relevant costs.
6. State transitions are realized, and the game proceeds to period $t + 1$.

Discussion of assumptions. While the AR(1) price forecast assumption is made for computational tractability purposes, I argue that that there is reason to believe that it approximates well firms' actual pricing forecasts. First, virtual machine prices stay relatively

[43] I set a convergence criterion that the maximum forecast error must be less than 0.1%. Because dynamic pricing games with switching frictions generally do not have unique equilibria (cf. Farrell and Klemperer (2007)), I begin each iteration using the empirical estimates for the AR(1) transition parameters during estimation. During the counterfactuals, I begin each outer iteration at the values obtained during estimation.

stable over time: the AR(1) persistence term for each supplier in the data is between 0.90 and 0.97. Second, prior research shows that firms have difficulty optimizing cloud costs (Brand et al. (2024)), with 65% (73%) of firms sampled by Flexera ranking "optimizing the existing use of cloud" as one of their primary cloud initiatives in 2019 (2020)(Flexera (2020, 2019)). Given the complexity of optimal VM provisioning under complete information about current pricing, it is unlikely firms form sophisticated beliefs about the future evolution of prices.

Second, the nonstationarity assumption is due to the fact that the cloud computing industry is rapidly growing during my sample period, with double-digit growth rates. This is captured by time-varying entry costs, κ_t^E, which represent the increased availability of software tools, such as Docker or Kubernetes, that ease the transition of workloads to the cloud and of cloud specialists. In addition, the average VM quality is rising over time due to a rapid rate of semiconductor innovation (cf. Lashkari et al. (2024)).

5 Estimation

To recover structural parameters, I estimate the model in two steps. First, I use a control function approach to estimate firms' production functions with endogenous technology choice (cf. de Loecker (2013); Ackerberg et al. (2015)), whose parameters are recovered using GMM. Firms' input choices reveal the output elasticities of inputs as well as the productivity gains from using cloud computing. Second, given production function coefficients, I estimate my remaining cost parameters using a method of simulated moments (McFadden (1989); Pakes and Pollard (1989)). Firms' investments and bundle choices recover entry and switching costs as well as brand specific maintenance cost shifters. While many parameters are jointly identified by the full set of moments in each stage, Table 4 lists parameters and provides brief intuition for the primary associated sources of identification for each. I report the results in Section 6, and provide details in Appendix C.2.

5.1 Estimation: Production function

To estimate the productivity gains from cloud computing, I first estimate a production function for each industry. While my sample size is sufficient to estimate production at the NACE 2-digit or even 4-digit level, this greatly increases the consumer type space and therefore computational burden of solving the dynamic model. Therefore, I follow Peters et al. (2017) and divide the manufacturing sector into high-technology and low-technology manufacturing. I define high-tech manufacturing as any 4-digit NACE code for which R&D expenditures are above 5% of gross value added, where values are taken from (Galindo-

Table 4: List of parameters and identification sources

Parameter	Source of identification
Production function parameters (GMM)	
$\boldsymbol{\beta}_x^i$ (Output elasticities of inputs)	Within-industry output and input covariation
$\boldsymbol{\gamma}_{0:3}^i$ (Productivity persistence)	Within-firm autocorrelation in investment
γ_4^i (Impact of cloud adoption on productivity)	Within-firm covariation of investment and cloud adoption
γ_5^i (Impact of COVID-19 on productivity)	Within-firm covariation of investment and 2020 year indicator
Bundle cost parameters (MSM)	
κ_t^E (Cloud adoption costs)	Annual share of cloud users
κ^S (Switching cost)	Autocorrelation of bundle choices among cloud users
α_0 (Maintenance cost intercept)	Exit rate from cloud computing among users
α_1 (Virtual machine price cost shifter)	Supplier price variation among new entrants and Bertrand conduct assumptions
α_2 (Multihoming cost shifter)	Share of cloud users that multihome
α_3 (Quality cost shifter)	Covariance of average quality and bundle choice among new entrants
α_4 (Azure-MS server cost shifter)	Covariance of Azure and Microsoft server usage among new entrants
ξ_b (Vertical differentiation)	Market shares of suppliers among new entrants
σ^E (Variance of entry costs)	Size premium of new cloud entrants
σ^b (Variance of maintenance costs)	Size premium of multihoming firms among cloud users
Investment parameters (MSM)	
C_1^K (Convex adjustment cost)	Industry level average investment rates
σ^K (Variance of investment costs)	Industry level variation in investment rates
Conduct parameter (MSM)	
ζ (Supplier conduct)	Supplier pricing vs. implied markups
Remaining parameters (calibrated)	
β (Discount rate)	Calibrated from Bloom (2009)
C_j^{vm} (VM marginal cost)	Calibrated using server depr. + electricity costs
C^{egress} (egress marginal cost)	Calibrated from Boudreau (2022)

This table summarizes the steps of my estimation procedure, detailed below. Each step lists the parameter, definition, and source of identification. For the first step, production function parameters are identified via Ackerberg et al. (2015). For the second step, parameters are estimated via method of simulated moments (MSM). Finally, several parameters are calibrated from prior academic or industry sources. An informal description of target moments are listed in column 2.

Rueda and Verger (2016)). Since services cover a large fraction of cloud computing demand, I include the sector as a third industry.[44]

Second, since I do not observe cloud computing usage at the intensive margin in my data, to avoid measurement error, I follow Peters et al. (2017) and assume technology choice is an ordinal variable (cf. Mairesse et al. (2005)).[45]

To estimate the production function, I follow the approach from Ackerberg et al. (2015) and use non-IT investment as a proxy variable to control for the simultaneity bias of productivity and investment choices. To construct non-IT investment, I first use data from Orbis to construct firms' total capital investment. My sample contains depreciation values for 62.9% of observations. For the remaining firms, I use firms' time series of capital stocks, along with industry-year level depreciation rates at the 4-digit NACE code level, to compute investment. I construct these depreciation rates from the entire sample of Orbis firms as $\frac{\text{Depreciation}_{ft}}{K_{ft}}$ and then take the mean. I then apply the perpetual inventory method to construct annual investment flows. Details are in Appendix A.5.

Three comments are in order about identification in my specific setting. First, to satisfy the monotonicity assumption of investment in productivity, I subtract PC, server, and storage investment from total investment to recover investment in non-IT capital. This is necessary because, in France, cloud services are expensed as services costs, while on-premise computing is capitalized (Dos Santos and Ovsyannikova (2023)). Thus, a firm may reduce on-premise computing investment in response to cloud adoption even as productivity rises. If computing is a relatively large share of a firm's capital stock, this could violate the assumption that I_{ft} is monotonically increasing in Ω_{ft}. By restricting my proxy variable to be non-IT investment, I increase the likelihood that monotonicity obtains, resulting in the following assumption:[46]

Assumption 1. *Non-IT investment, $\tilde{I}_{ft}$, increases monotonically with firm productivity, Ω_{ft}.*

[44] Following standard practice, I winsorize each sector at the 1st and 99th percentile of firms' capital size (cf. Rubens et al. (2024)). In addition, this allows me to exclude from the analysis large firms who are likely to bilaterally bargain with suppliers over price and other contract terms.

[45] One could use model assumptions to back out IaaS spending using firms' cost minimization equation and current on-premise computing usage. However, this involves making strong assumptions about computing capital utilization rates, which are known to very widely by type of computer as well as firm and industry (Brynjolfsson and Hitt (1996); Brand et al. (2024)).

[46] There are potentially two issues with this proxy variable. First, using capital investment as a proxy creates a selection on firms with positive investment. In my sample, 80.3% of high-tech manufacturing firms, 77.3% of low-tech manufacturing firms, and 69.9% of services firms have positive investment rates. Second, Aberdeen estimates IT investment expenditures using data on IT stocks collected from interviews. Given the opacity of their modeling procedure, this may introduce measurement error. As a robustness check, I use total investment as my proxy variable. Results are presented in Appendix D.1.4 and are largely consistent with my primary specification.

Second, I need to control for the endogeneity of firms' cloud computing choices in productivity. It is likely, for example, that more productive firms are more likely to select into cloud usage. As well, finding a suitable instrument is difficult, due to most instruments likely also affecting other IT investments.[47] To control for this selection effect, I exploit the timing assumptions of my productivity process. In particular, because I assume that firms make cloud adoption decisions one period in advance, it is orthogonal to firms' current period productivity shock. I follow de Loecker (2013) and exploit the following moment condition to identify γ_4:[48]

$$\mathbb{E}[b_{ft} v_{ft}] = 0 \tag{24}$$

Likewise, I assume that the COVID-19 shock is drawn after firms make their flexible input choices in $t = 2020$.

Third, I control for the impact of input and output quality on productivity.[49] I make the assumption that firm input and output prices are increasing in (unobserved) product quality. In the baseline specification, I use firm level wages, P^l_{ft}, to proxy for unobserved worker quality. I observe the total wage bill for a majority of firms. For the remaining firms, I impute the wage at the year-region-industry level.[50] To estimate the input price control function $H(p^l_{ft})$, I estimate a third order polynomial in the first stage of my control function approach:

$$\tilde{\omega}_{ft}(\boldsymbol{\beta}, \boldsymbol{\vartheta}) = \hat{\phi}_{ft} - f(\mathbf{x}_{ft}; \boldsymbol{\beta}) - H(p^l_{ft}; \boldsymbol{\vartheta}) \tag{25}$$

where $\boldsymbol{\beta}, \boldsymbol{\vartheta}$ are the production function and input control function parameters, respectively; $\hat{\phi}(\cdot)$ is predicted output; and $\mathbf{x}_{ft}$ is a vector of inputs, including b_{ft}. Finally, to target the production function parameters $\boldsymbol{\beta}$, I instrument flexible inputs with their lags. Formally, the relevant moment conditions are

$$\mathbb{E}[v_{ft}(\boldsymbol{\beta}, \boldsymbol{\vartheta}) \mathbf{I}_{ft}] = 0 \tag{26}$$

[47] Though see DeStefano et al. (2023) for an exception.

[48] It may be that firms, especially smaller firms, upgrade to cloud computing within a year of making an adoption decision. As a robustness check, I instead assume that $\mathbb{E}[b_{ft-1} v_{ft}] = 0$ and productivity gains are only realized one period after adoption is observed. As a second check, I examine whether firms experience gains from learning-by-doing (Brand et al. (2024)) and use the number of years using cloud computing as my independent variable.

[49] See De Loecker et al. (2016) for a discussion.

[50] To reduce the size of the state space in my dynamic model, I only include wages as a proxy for input quality. In Appendix D.1.5, I include the additional quality proxies from De Loecker et al. (2016), namely export status and firm market share at the industry-region-year level. Results are broadly similar to my primary specification.

where $\mathbf{I}_{ft}$ contains current capital, lagged materials and labor, and their higher order and interaction terms, as well as lagged wages.

5.2 Estimation: Cost parameters

Given estimates of the production function parameters, I estimate the remaining parameters via the method of simulated moments (MSM). While the full moment set jointly identifies the full parameter set, I provide intuition for the moments chosen primarily to identify each parameter. The common component of the entry cost κ_t^E is primarily identified from the annual fraction of firms that adopt cloud computing. The variance of entry cost shocks, σ^E, is identified from the size premium of adopters over non-adopters (Arcidiacono and Miller (2011)). To see the intuition for this, note that, given the assumption on the form of adoption cost shocks, the log odds ratio that a firm adopts some cloud computing bundle is given by

$$\log\left(\frac{\mathbb{P}_t[b_{ft+1}\neq\varnothing|b_{ft}\neq\varnothing]}{1-\mathbb{P}_t[b_{ft+1}\neq\varnothing|b_{ft}\neq\varnothing]}\right)=-\frac{\kappa_t^E}{\sigma^E}+\frac{1}{\sigma^E}\cdot\left(\frac{\beta}{D_{ft}}\max_{b_{ft+1}}\mathbb{E}_tV_{t+1}(b_{ft+1},\tilde{s}_{ft+1})-\frac{\beta}{D_{ft}}\mathbb{E}_tV_{t+1}(\varnothing,\tilde{s}_{ft+1})\right) \tag{27}$$

As long as profits and the data stock D_{ft} do not perfectly scale with size, then σ^E is identified by the size premium of adopters. Next, I identify switching costs κ^S using the autocorrelation of cloud bundle choices among cloud users. To identify the level of maintenance costs α_0, I use the exit rate from cloud computing among cloud using firms. Likewise, the additional multihoming cost α_2 is identified by the rate of multihoming among cloud users. For the remaining moments, I focus on new cloud adopters to separate firm preference heterogeneity from switching costs (cf. Heckman et al. (1981)). To target α_1, I use the covariance of price and bundle choice among new cloud entrants and among new entrants without Microsoft server products. In addition, my Bertrand conduct assumption for suppliers provides an additional source of identification for the price parameter, since I use average supplier prices across datacenter-years as additional moments. To identify α_3, I use the covariance of choosing a bundle with Azure and using Microsoft server products. To target unobserved vertical quality, $\xi_{\text{Microsoft}}, \xi_{\text{Amazon}}$, I use the market shares of Amazon and Microsoft among new entrants.[51] Finally, to identify the variance in maintenance cost shocks σ^b, I use the size premium of multihoming firms among cloud users, following analogous reasoning to the identification of σ^E.

To identify the convex cost parameter and variance of capital adjustment costs, I use the

[51] The author's interview with a chief cloud architect who has worked in the finance and healthcare industries suggests that Google has the lowest quality cloud platform among the major providers, so to separately identify α_0 and ξ_j, I normalize ξ_{Google} to 0.

mean and variance of the investment-to-capital-stock ratio, I_{ft}/K_{ft}, respectively. Finally, given calibrated marginal costs, the mean price of each supplier across datacenter-years identifies the harvest/invest conduct parameter.

To estimate these parameters $\boldsymbol{\theta} = \{\kappa_0^E, \kappa_1^E, \kappa^S, \boldsymbol{\alpha}_{0:4}, \boldsymbol{\xi}_j, C_1^K, \sigma^K, \sigma^E, \sigma^b, \zeta\}$, I use a full solution approach. At each guess of the parameter vector $\boldsymbol{\theta}$, I solve the an outer and inner loop. For each iteration of the outer loop, firms guess an AR(1) price transition kernel. In the inner loop, I solve for firm investment and technology choice policy functions at the steady state period T sufficiently far in the future, when entry costs are low enough such that cloud computing has fully diffused.[52] I set $T = 2045$. I then use backwards induction to compute downstream firms' investment and computing bundle choices and upstream firms' pricing strategies. I then draw an initial distribution of states, $F(s_0)$, simulate the model forward and run AR(1) "regressions" over the time series of prices for each supplier, comparing the guessed and true pricing kernels. I repeat until convergence, at which point I compare the observed and simulated moments. Given that the MME outer loop often fails to converge far from the global optimum, to make this step computationally feasible, I follow Gortmaker (2024) and first use Bayesian optimization, a derivative-free global approach, to obtain an initial estimate. I then employ a trust-region refinement with numeric derivatives to obtain the final estimate. To estimate the optimal model parameters, I first aggregate the industry-level predictions to the country level by taking a weighted average, as in Asker et al. (2014):

$$\boldsymbol{\Psi}(\boldsymbol{\theta}) = \frac{1}{\sum_i N_i} \sum_i N_i \boldsymbol{\Psi}_i(\boldsymbol{\theta}) \tag{28}$$

where N_i denotes the number of firms in industry i. To target each parameter, I simulate statistics, which I average over firm-industry-years and form simulated statistics. Then the moments are matched to the data moments $\hat{\boldsymbol{\Psi}}$. To reduce the effect of outlier moments, I estimate parameters by minimizing the criterion function given by the usual quadratic form, with weighting matrix $\mathbf{W}$:

$$\mathcal{Q}(\boldsymbol{\theta}) = (\boldsymbol{\Psi}(\boldsymbol{\theta}) - \hat{\boldsymbol{\Psi}})' \mathbf{W} (\boldsymbol{\Psi}(\boldsymbol{\theta}) - \hat{\boldsymbol{\Psi}})$$

where I set $\mathbf{W}$ is a diagonal weight matrix with values set such that all moments are roughly on the same scale. In particular, moments with a value above 1 get a weight of 0.1, while all other moments get a weight of 1. Standard errors are calculated using moment vectors from the same bootstrapped samples. For each industry, I discretize Ω_i to 20 grid points and K_i to 20 grid points, evenly spaced between the 1st and 99th percentile of the

[52] At the steady state, I additional assume that cloud supplier prices and qualities are held fixed at, respectively, $P_{jT}, \log(q_{jT})$.

data.[53] Finally, I assume the COVID-19 productivity shock is unanticipated so I simulate the model twice for each parameter guess - once for 2017-2019 without the shock and once for 2020-T with it. Further details are in Appendix C.2, and computational details are presented in Appendix C.3.

6 Estimation results

To study the impact of cloud computing adoption on firm productivity, I present the results of my production function estimation. Using those results, I can then recover the cost parameters. I find small but significant effects of cloud adoption in all industries, with larger productivity gains in manufacturing and among firms that multihome. I find that adopting Microsoft or Amazon leads to lower costs than adopting Google; that firms are fairly sensitive to price; and that there is a strong complementarity between Microsoft's cloud and on-premise server products. In addition, I compare my estimates to prior research and industry case studies on the benefits and/or costs of adopting cloud computing.

6.1 Production function parameters

Production function coefficients are presented in Table 5. Results are significant for most coefficients and reasonable, with services having a lower capital share than manufacturing.

Table 5: Production function estimates

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Labor (β_l)	0.114 (0.094)	0.286 (0.048)***	0.267 (0.038)***
Materials (β_m)	0.617 (0.051)***	0.480 (0.079)***	0.706 (0.310)*
Capital (β_k)	0.084 (0.027)***	0.085 (0.017)***	0.034 (0.015)*
Wage (ϑ_w)	0.139 (0.067)*	0.285 (0.046)***	0.148 (0.112)
Observations	14,893	37,948	153,230
Firms	4,877	12,791	70,321

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) wage control for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

Next, I present coefficients for the productivity process in Table 6. The coefficient on cloud computing adoption is small yet significant across all specifications. Thus, services

[53] Prior versions of the project estimated the model with 30-point grids for Ω_i and K_i, producing similar parameter estimates. I choose the smaller grids for computational tractability.

Table 6: Productivity estimates

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Productivity intercept (γ_0)	0.160 (0.022)***	0.154 (0.018)***	0.140 (0.003)***
Linear AR(1) term (γ_1)	0.719 (0.050)***	0.686 (0.044)***	0.499 (0.011)***
Quadratic AR(1) term (γ_2)	0.167 (0.036)***	0.229 (0.036)***	0.509 (0.013)***
Cubic AR(1) term (γ_3)	-0.034 (0.009)***	-0.057 (0.009)***	-0.167 (0.004)***
Cloud adoption (γ_4)	0.005 (0.001)***	0.004 (0.001)***	0.002 (0.000)***
COVID-19 shock (γ_5)	-0.049 (0.001)***	-0.054 (0.001)***	-0.034 (0.001)***
Observations	14,893	37,948	153,230
Firms	4,877	12,791	70,321

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) wage control for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

firms that adopt cloud computing experience productivity gains of 0.2%-0.6%, depending on number of providers used. The equivalent gains are 0.4%-1.2% in low-tech manufacturing and 0.5%-1.5% in high-tech manufacturing. Figure 6 plots the steady state sector-level productivity distributions for firms that adopt a single cloud provider vs non-adopters. Both manufacturing sectors show a clear benefit of cloud computing. The benefit to firms in the services sector is less noticeable, both because the distribution of productivity varies less and because of the smaller coefficient.[54] Finally, Figure 7 displays the median difference in productivity paths for a firm in each sector, comparing a non-adopter with a firm that has adopted one cloud provider.

Are these effects reasonable? Most studies of the impact of cloud computing tend to merge IaaS with PaaS and/or SaaS products (e.g. DeStefano et al. (2023)). The two exceptions are Jin and Bai (2022), who use the most similar definition to mine in Table A3.[55] They find a productivity increase of 0.47% per cloud-related job posting in the prior quarter. Using the mean value of this variable, 6.441, this comes to productivity gains of 3.03%. This effect may be due to the fact that their results focus on public US firms, which experience higher productivity gains from IT adoption than non-US firms due to better management practices (Bloom et al. (2012)). As another exercise, I compute the gains from adopting Microsoft Azure IaaS in the case study reported in North et al. (2019). Including cost savings, and ignoring equilibrium effects of non-cloud inputs, their results imply a productivity gain on average of 0.17% for the first 3 years when including cost savings and a gain of 0.05%

[54] Appendix D.2 displays the results for firms adopting all 3 cloud suppliers.

[55] This measure includes IaaS, PaaS, and SaaS computing and storage-related postings but ignores distributed computing and cloud server postings.

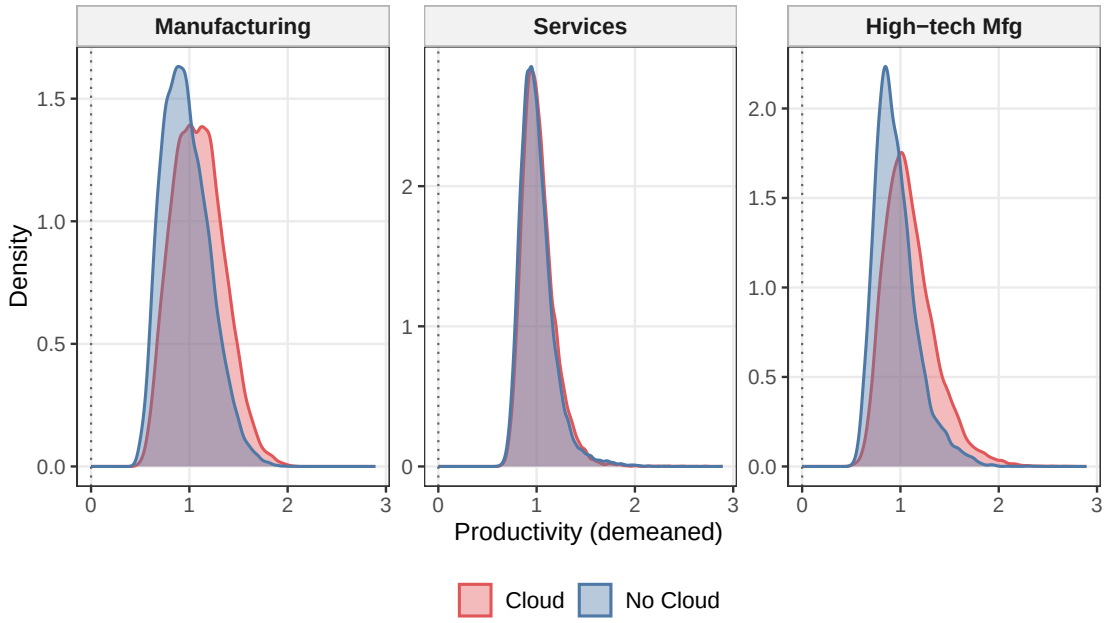


Figure 6: Steady state productivity distribution

Note: Simulated steady state productivity distributions for firms that adopt 1 cloud supplier vs those that have not adopted any suppliers.

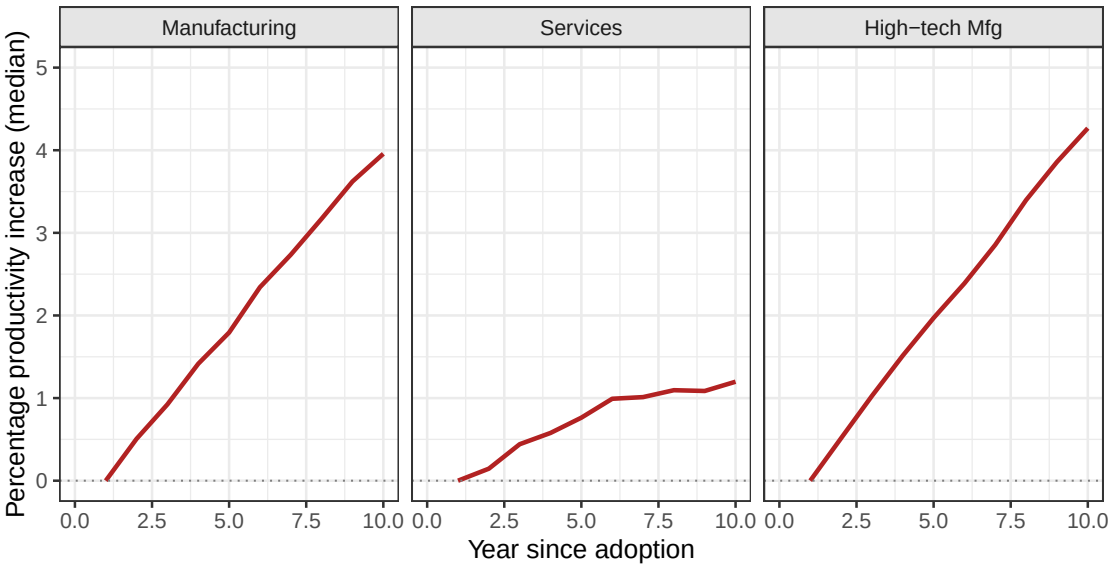


Figure 7: Median productivity difference between adopters and non-adopters.

Note: Difference in median productivity for adopters and non-adopters, from 10,000 simulations over 10 periods.

when only including increases in revenue.[56] While empirical evidence remains scarce, it is heartening that my estimates fall between these two.

6.2 Dynamic parameters

Using my production function estimates, I then estimate the remaining parameters via MSM. Before displaying the results, I estimate the relationship between the capital stock, K_{ft}, and its data stock, D_{ft}. Because Aberdeen only contains data on on-premise storage investments, I use data from 2015-2016 before cloud computing was widely adopted. In order to reduce the state space, I recover D_{ft} as a function of a firm's capital stock K_{ft} using the following equation:[57]

$$\log(D_{ft}) = \upsilon_{0i} + \upsilon_{1i} \log(K_{ft}) + \Xi_{\iota} + \tau_t + \epsilon^D_{ft} \tag{29}$$

where D_{ft} is data stock in petabytes[58], Ξ_{ι} are 4-digit NACE code industry fixed effects, τ_t are time fixed effects, and $\epsilon^D_{ft} \sim \mathcal{N}(0, \sigma^D)$. To convert to PB, I first convert to USD using the average annual exchange rate. I scrape HardDriveBenchmark.net (nd) for HDD prices and volumes back to 2018 and compute the annual average cost in \$/PB.[59] Then, I use the US BLS price index for disk storage to obtain the price per PB back to 2015 (McCallum and of Labor Statistics (2024)).[60] Coefficients for this regression are presented in Table 7, showing a positive and significant relationship between data investments and total investments.

Given these estimates, Table 8 reports the estimates of the dynamic choice parameters. Entry costs are falling rapidly; for a firm with a capital stock of €$4.8M$ (€$0.5M$), the cost falls from €$7.63M$ (€$5.48M$) in 2017 to €0.54 (€0.39) by 2021, consistent with the rapid rise in adoption observed in the data. Switching costs are relatively high, at around €2.98 per GB of data stored, but a majority in the variation in supplier choice is coming from price and feature parameters. For example, a doubling of the price per core-hour increases maintenance costs by €18.2 per GB of data in the cloud. The complementarity of Microsoft Azure and on-premise Microsoft server products is especially large to rationalize the high

[56] For full details of the comparison, see Appendix D.4.

[57] For a derivation, see Appendix B.4.

[58] 1 PB = 1 million GB.

[59] I use hard drives rather than solid state drives, as prior research has shown that even by 2016, most storage volume sold was on HDDs rather than SSDs (Igami and Uetake (2020)). Recent industry reports corroborate that this trend continued into more recent years (TRENDFOCUS, Inc. (2019)).

[60] The narrowest consumer price index in Eurostat that contains storage is HICP 09.1.3 ("Peripheral equipment"), which bundles goods with a rapid price decline such as storage with others like monitors and keyboards. Process innovation in storage media is the largest driver of annual price changes for storage, rather than cross-country currency fluctuations between the US and Eurozone. Therefore, I use the more precise BLS index.

Table 7: Elasticity of data stock with respect to capital stock

Variable	High-tech manufacturing	Low-tech manufacturing	Services
υ_0	-1.209 (1.059)	-0.364 (0.291)	-0.409 (0.187)*
υ_1	0.427 (0.026)***	0.365 (0.024)***	0.210 (0.014)***
Observations	1,272	1,810	3,166
Firms	769	1,164	2,062
Industry FE	✓	✓	✓
Year FE	✓	✓	✓
R^2	0.2885	0.2800	0.1982
Adj. R^2	0.2513	0.2261	0.1560

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions estimated using OLS with NACE4 and year fixed effects.

covariance between Microsoft products observed in the data. The term on log quality is relatively small, consistent with prior research that firms take many years after adoption to optimize their workloads across types of processors (Brand et al. (2024)). Finally, the vertical quality terns ξ_j are both positive and significant, consistent with the fact that AWS and Microsoft Azure are relatively more developed platforms than Google Cloud Platform.

Model fit is reported in Figure 8 for entry rates. As shown, the model fits entry dynamics relatively well, if the adoption rate is a bit low in some years. In Appendix C.2, Table C.2 presents how well the model fits my full set of chosen moments.

Using this data, I can estimate firms' cost of using cloud computing. Table 9 presents estimates of the costs of using various cloud providers in the manufacturing sector, for firms at various sizes in the size distribution used in the simulation.[61] As expected, Microsoft and Amazon are lower priced than Google, largely due to high vertical quality terms that rationalize their higher adoption rates relative to Google. How do these numbers compare to real-world estimates? While representative data are difficult to obtain, an analysis by North et al. (2019) estimates that a firm with revenue of €3.6 billion saw 3-year entry and maintenance costs of using Microsoft Azure of €4.9 million. Among firms of similar size, I find that the corresponding costs are €0.44-5.97 million in 2019. Likewise, Forrester estimates of the gain in profit are 0.51% over 3 years, while I find a corresponding estimate of 1.57%-3.75%. While these numbers are not directly comparable (for example, Forrester samples a US public school district), it is reassuring that they are on similar orders of magnitude, especially on the cost estimates.[62] Likewise, Autorité de la concurrence (2023)

[61] In Appendix D.3 I display the corresponding estimates for the services and high-tech manufacturing sectors.
[62] For the full back of the envelope calculation, as well as results for other sectors, see Appendix D.4.

Table 8: Cost Parameter Estimates

Parameter	Estimate	95% CI
Entry Costs		
Constant (κ_0^E)	3.5916	[3.1859, 3.6489]
Slope (κ_1^E)	-0.6638	[-0.7241, -0.5962]
Switching cost		
Switching cost (κ^S)	2.9845	[2.7619, 3.1809]
Maintenance Costs		
Baseline (α_0)	2.4330	[1.6542, 2.4941]
Price (α_1)	18.2471	[18.2178, 18.5876]
Multihoming cost (α_2)	0.8132	[0.6730 0.8557]
VM log quality (α_3)	-0.3831	[0.2032, 0.4861]
Azure-MS server complementarity (α_4)	-4.0511	[2.4717, 4.1718]
Brand Effects		
Amazon brand effect (ξ_{AWS})	-0.5710	[0.3149, 0.6027]
Microsoft brand effect (ξ_{MS})	-1.4477	[1.0437, 1.5129]
Capital Costs		
Convex adjustment cost (C_1^K)	2.1044	[1.2852, 2.2293]
Cost Heterogeneity		
Entry cost shock variance (σ^E)	1.0594	[0.5190, 1.1048]
Maintenance cost shock variance (σ^b)	0.9742	[0.8591, 1.0227]
Investment cost shock variance (σ^K)	2.2686	[2.2500, 2.3001]
Supplier Conduct		
Supplier conduct (ζ^S)	0.1003	[0.1000, 0.1005]

Notes: Estimates are recovered using MSM. 95% confidence intervals are percentiles from 80 block bootstrapped samples, which reflect sampling uncertainty from second-stage moments but not production function parameters.

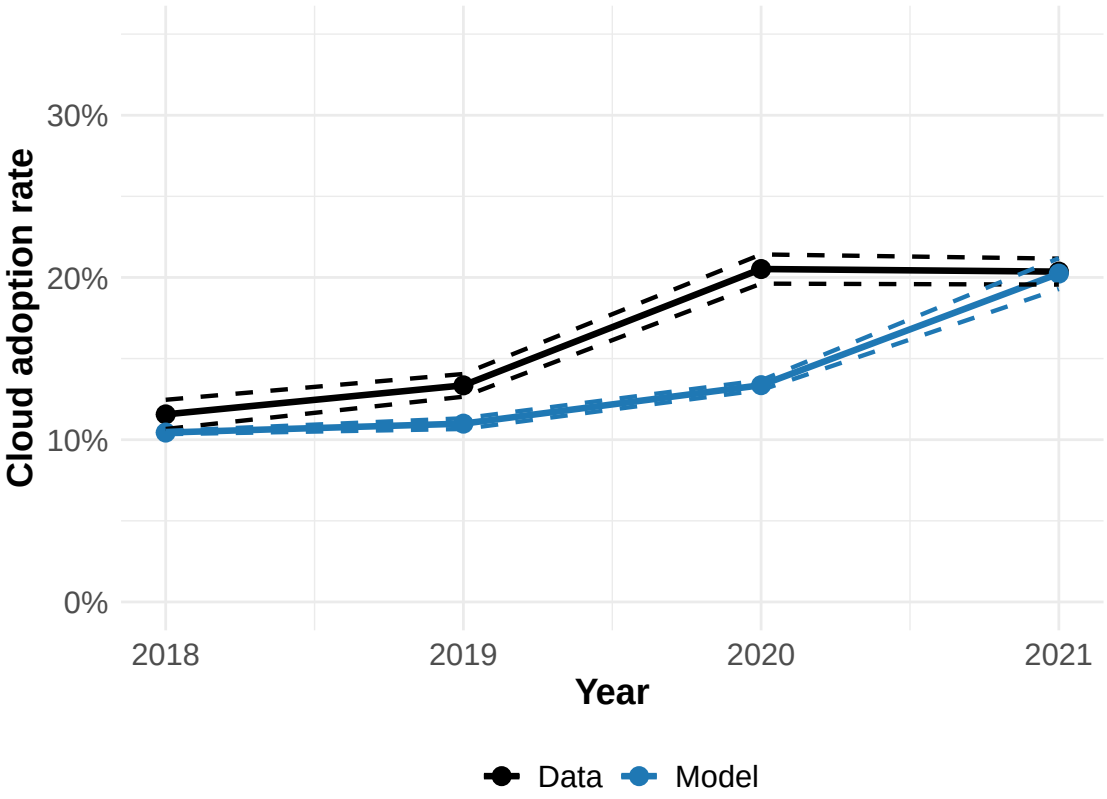


Figure 8: Actual vs simulated cloud adoption rates
Dashed lines represent 95% CI, which account for sampling and simulation uncertainty for the simulated moments and sampling uncertainty for the data moments.

finds that in 2022, the median cost to adopt cloud computing among survey respondents was over €450,000, quite similar to my 2021 estimate, while the mean cost was €6.6 million.

Table 9: Manufacturing sector cloud cost distribution

Maintenance costs					
Supplier	Minimum	10th Percentile	Median	90th Percentile	Maximum
Microsoft	0.1182	0.1548	0.2015	0.2805	0.3941
Amazon	0.1842	0.2411	0.3138	0.4369	0.6139
Google	0.2562	0.3353	0.4365	0.6076	0.8539
Entry costs					
Year	Minimum	10th Percentile	Median	90th Percentile	Maximum
2017	3.2181	4.2128	5.4836	7.6334	10.7272
2021	0.2262	0.2961	0.3854	0.5365	0.7540

Note: Costs computed for firms used in the simulation, in millions of real euros. Supplier prices and qualities are computed using 2017 values. I assume firms are not using Microsoft server products.

7 Counterfactual policies

To study the welfare of competition policies for cloud computing, I use my model to conduct counterfactual scenarios. First, to directly quantify the impacts of input market power, I consider two scenarios. To simulate the impact of increased competition, I consider the impact of an additional entrant in 2017 on prices and welfare. I assume that this entrant has the same unobserved vertical quality as Google. Then, I simulate a merger between Microsoft Azure and Google Cloud Platform, conservatively holding fixed egress fees and disallowing additional merger synergies. Second, I simulate the impact of the two relevant competition policies. I first study the impact of a ban on data egress fees implemented in the 2024 EU Data Act and studied by other regulators (e.g. Korea Fair Trade Commission (2022)). Next, I consider the case in which firms using Microsoft Server products see gains from using these products on Google or Amazon cloud, as would be the case if a regulator mandated price parity or higher feature compatibility for using Microsoft products on rival cloud platforms. For each scenario, I compute the average annual changes in producer and consumer surplus, prices, and adoption from the baseline. I also study how each evolve over the industry's expansion path.[63]

[63] Appendix E.1 presents a stylized model to build intuition about how the two competition policies can affect equilbrium supplier pricing.

7.1 Computation

Recalling I make the simplifying assumption that prices in the steady state are fixed at time T prices, to minimize the impact of bias, I compute welfare along $t = 2017, ..., \tilde{T}$, where $\tilde{T} = 2031$. Per period welfare can be written as

$$\mathbb{W}_t = \underbrace{\int N_t(s_{ft})[\pi_{ft}(s_{ft}) - \kappa_t(D_{ft}) - C(I_{ft}, K_{ft})]dF(s_{ft})}_{\text{Consumer surplus}} + \underbrace{\sum_j \pi^S_{jt}(\mathbf{s}_t)}_{\text{Producer surplus}} \tag{30}$$

where $N_t(s_{ft})$ is the mass of firms at state s_{ft} and $dF(s_{ft})$ is the joint distribution of states. Then, total social welfare can be written as

$$\mathcal{W} = \sum_{t=2017}^{\tilde{T}} \beta^t \mathbb{W}_t \tag{31}$$

To solve the counterfactual, I initialize the model using the observed 2017 distributions over states. Next, I assume that the policy arrives as an unexpected shock after firms have drawn their initial states, affecting all states in period $t = 2018, ..., T$. For the counterfactual scenarios without price increases, I simply re-solve each period's value function fixed point holding prices fixed. For the counterfactual with endogenous price updates, I unfreeze the MME loop and re-solve for prices and downstream firms' beliefs. Finally, counterfactuals can change suppliers' invest-harvest incentives. In the appendix, I provide a a quasi-Newton fixed point to solve for ζ^{CF}. I cannot guarantee equilibrium uniqueness, so I begin each iteration at the equilibrium such that counterfactual equilibria are "near" the factual.[64]

7.2 Impacts of market concentration

To study the direct impacts of market concentration on cloud adoption and industry welfare, I simulate the entry of a new cloud provider with the same characteristics as Google.[65] To isolate the price effects and avoid extrapolating my productivity results out of sample, I assume no additional productivity gains from adopting a fourth supplier. Adopting a fourth supplier incurs an additional multihoming cost, however. Next, I simulate a merger between Microsoft Azure and Google Cloud Platform, such that the upstream market becomes a duopoly.[66] Microsoft and Google are natural candidates for such a simulation, given that

[64] See Appendix E.2 for details of the algorithm and Appendix E.4 for results holding ζ fixed.

[65] That is, $\xi_{Entrant} = 0$ and $q_{Entrant,t} = q_{Google,t}$.

[66] The closest such example in the industry is IBM's 2019 acquisition of Red Hat, a firm that provides open-source cloud-adjacent software and hardware, such as containerization, hybrid cloud infrastructure, and the Red Hat Enterprise Linux OS.

Amazon worldwide remains the market leader. In addition, since Google is the market laggard, allowing these firms to merge could reduce double marginalization among multihomers that use Google as a second or third supplier. Given that this scenario has not occurred in the data, I make the conservative assumption that the merger only results in joint pricing decisions by the merged firm, without additional synergies that result either in reduced egress costs or lower marginal costs.[67]

7.3 Ban on data egress fees

The first policy is a ban on data egress fees by cloud providers. The rationale behind such a ban is as follows. Egress fees act as a form of switching cost to consumers, locking them into long-run relationships with suppliers, even when better alternatives are available. Even with a rapidly-expanding consumer base, the high fixed costs of entry on the supply side means that the market remains concentrated. Theoretical analysis of an egress fee ban are mixed, with welfare results dependent on the demand elasticity (Do and Miklós-Thal (2025); Gans et al. (2023)).[68]

To simulate the counterfactual scenario, I assume that, beginning in 2017, cloud computing suppliers are banned from charging egress fees, and simulate sectors forward until T, computing welfare until $\tilde{T}$.[69] If a firm switches from a single supplier to another or exits the cloud computing market entirely, I assume they pay the full egress fee schedule listed under "egress to the public internet" or the equivalent. Otherwise, I do not observe the fraction of data a firm stores in each supplier when multihoming, which is endogenous to egress fees (Competition and Markets Authority (2025c)). I assume that when firms multihome, they keep an equal fraction of data on all providers. Therefore, if a firm switches from $b_t = \{\text{Amazon}\}$ to $b_{t+1} = \{\text{Amazon, Microsoft}\}$, it will transfer 50% of its year t data stock. For each firm, I subtract this cost from κ_S. In addition, firms may move a fraction of their data between suppliers in each period when multihoming. For example, a firm may need to move customer data from its website, hosted on Microsoft Azure, to be analyzed using machine learning tools, hosted on Google Cloud Platform. Recall that I assume that firms move a fraction $\lambda = 0.15$ of data between providers in each period t. Finally, since egress

[67] The lack of a reduction in egress costs is motivated by the fact that all major cloud providers charge nonzero egress fees for moving data across *their own* datacenter locations, such that even within-supplier datacenter switching is rare (cf. Demirer et al. (2024)).

[68] Prior research on number portability in telecommunications, however, has tended to find that it reduced prices and increased consumer surplus (Park (2011); Viard (2007)).

[69] In Appendix E, I consider the scenario in which cloud firms do not update prices in response to the egress fee ban. This follows medium term pricing dynamics: as of June 2025, none of the major cloud providers have raised IaaS prices in response to the removal of egress fees in January and March of 2024, respectively (Competition and Markets Authority (2025g), 44).

fees are in GB/month, I follow Competition and Markets Authority (2024) and assume that multihoming firms move data across platforms uniformly during each year.[70]

7.4 Increased Microsoft server interoperability

The second policy I consider is a higher level of adoption benefits for firms using Microsoft server products on non-Microsoft cloud platforms. The UK Competition and Markets Authority has found that Microsoft charges more for Windows Server licenses for use on on rival cloud platforms than on Microsoft Azure, leading them to open an investigation into anticompetitive bundling. In addition, programs such as Azure Hybird Benefit, which allows consumers to host Azure services on excess local computing capacity for a reduced cost, further increase the complementarity between the two products. I capture the impact of a hypothetical mandate to increase price or feature parity for Microsoft server products across products through allowing firms not using Microsoft Azure to capture a portion $\varepsilon \in \{0.1, 0.2, 0.5\}$ of the benefit that Azure users do by reducing maintenance costs by $\varepsilon \cdot \alpha_4$ if they use Microsoft server products.[71] In this section, I present results for $\varepsilon = 0.5$, with additional results in Appendix E.

Such a policy has contradictory effects for Microsoft Azure. On the one hand, it lowers the price of cloud computing, increasing the number of adopters, including multihoming firms. On the other hand, such a policy is likely to reduce the market share of Microsoft among cloud users. Note that such a policy is likely to have second order effects as well - ensuring interoperability or price parity is likely to reduce Microsoft's incentives to innovate in its server products, and I do not consider endogenous adoption of and revenue from server products. These effects are left to future research.

7.5 Results

The results are displayed in Table 10.[72] First, panel A displays the results for adding an entrant and for the merger. In the entrant scenario, average annual welfare increases by €8.4 billion annually, or 2.82%. This is partly driven by new entry, which rises by 4.53% percentage points on average. However, multihoming greatly increases (by 18.57%), driven by a new low-cost entrant–to gain market share, the entrant prices at an average annual of

[70] The fraction of data moved monthly, then, is $\frac{\lambda D_{ft}}{12}$.

[71] I do not observe how much of the estimated complementarity from Microsoft Azure and Microsoft server products is due to cost, software interoperability, or existing supplier relationships. I therefore test the robustness of the counterfactual to multiple values of ε.

[72] In Appendix E, I test the robustness of my results holding the invest-harvest conduct parameter ζ fixed at the baseline $\hat{\zeta}$.

€0.032. The welfare gains to the entrant and to Microsoft, which sees an increase in profits of 18.31% due to multihoming, outweigh the losses to Amazon and, especially, Google. In contrast, further consolidation of the sector does not produce as large a loss in welfare. Consumer surplus falls by €837.5 million (0.28%) annually, almost an order of magnitude lower, even though prices rise by 25.9%. Cloud products are undifferentiated enough such that the gains from increased profit per VM sold are offset by a reduction in cloud adoption among downstream firms, leading to a small producer surplus loss of 0.51%.

Turning to the impact of the two competition policies in panel B, we see that both produce positive but relatively small welfare gains. The egress fee ban produces a relatively small overall increase in annual welfare - €192.6 million (0.06%) annually. While gains to consumers are €313.8 million (0.11%), this is offset by a fall in supplier profits of 0.91%. The loss of egress fees, which make up 1.13% of supplier revenue, are slightly offset by an increase in adoption.[73] Assessing the goals of the EU Data Act and other regulators–to remove barriers to adoption, switching, and multihoming–we see that there is negligible impact on adoption (only 0.36% increase) and a slightly larger impact on multihoming (increase of 1.53%). Switching likewise rises, but by only 3.54% on average. The increase in Microsoft server compelementarity produces a slightly larger increase in consumer surplus (0.16%) and a small rise in producer surplus, driven by increased adoption of Amazon and Google at the expense of Microsoft. Overall welfare rises by an average €512.7 million (0.54%) annually.[74]

What drives these results? Figure 9a displays the average annual price of suppliers. Both scenarios fall in between the price changes driven by either entry or further consolidation. In particular, while the decline in prices from the egress fee ban is small at first, it increases over time as suppliers move from investing to harvesting. In contrast, if interoperability between Microsoft server software and rival cloud platforms rises, the market sees an immediate fall in supplier prices. I find that all suppliers lower their prices in response - Amazon and Google to compete with newly-interested Microsoft server users and Microsoft to retain its customer base. As the market matures, however, the magnitude of the fall in prices declines, as downstream firms' switching costs become the primary obstacle to competition. Turning to 9b, we see that adoption rises under both policies, though the difference is larger midway through the transition path: in the long run, the falling adoption cost will eventually drive nearly all firms to use cloud computing.

How does consumer surplus evolve then? As shown in Figure 10, consumer surplus rises

[73] My estimate of egress fees as a percentage of supplier revenue aligns with findings by the UK CMA, which found that egress fees made up $\leqslant 5\%$ of each supplier's UK cloud computing revenue during 2020-2023 (Competition and Markets Authority (2025e)).

[74] See Appendix E for results for $\varepsilon = 0.1, 0.2$.

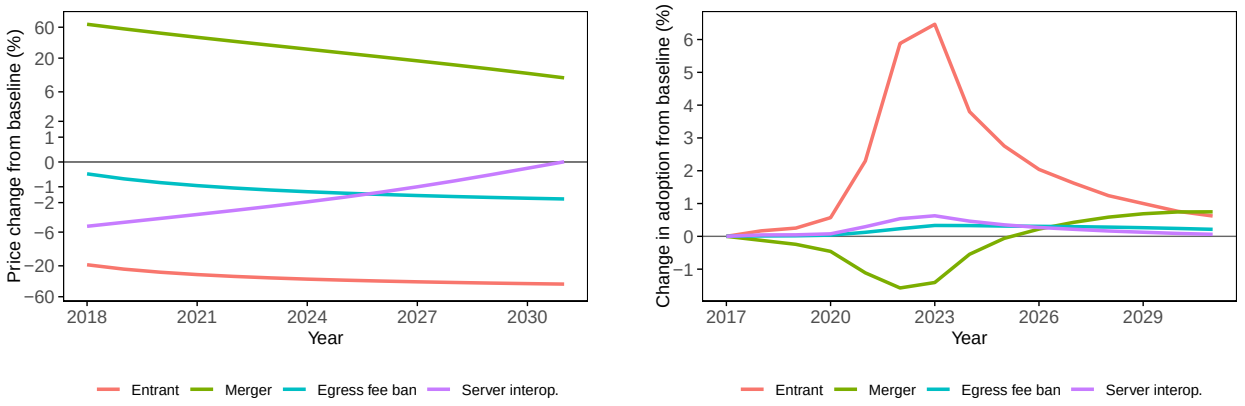
	ΔCS	ΔPS	ΔWelfare	ΔPrice	ΔAdoption	ΔMultihoming
A: Changes in market structure						
New entrant	6495.1	1940.1	8435.3	-0.023	1.97pp	5.60pp
	(2.24%)	(14.65%)	(2.82%)	(-30.68%)	(4.53%)	(18.57%)
Google-Microsoft merger	-837.5	-67.8	-905.3	0.020	-0.14pp	0.96pp
	(-0.28%)	(-0.51%)	(-0.29%)	(25.90%)	(-1.10%)	(3.25%)
B: Counterfactual policies						
Egress fee ban	313.8	-121.2	192.6	-0.001	0.20pp	0.48pp
	(0.11%)	(-0.91%)	(0.06%)	(-1.23%)	(0.36%)	(1.53%)
Server complementarity	489.4	27.8	517.2	-0.001	0.22pp	-0.28
	(0.16%)	(0.21%)	(0.16%)	(-1.66%)	(0.54%)	(-0.65%)

Table 10: Welfare effects of counterfactual exercises

Notes: Welfare changes relative to the estimated baseline over 2017-2031, in millions of real 2015 euros, averaged over 80 block bootstrap samples. Adoption the average annual share of firms using cloud computing. Multihoming is the average annual percentage of firms using multiple cloud suppliers. Price is the unweighted average annual price among suppliers $j \in \mathcal{J}$.

initially under the policy shock, as downstream firms benefit from lower prices.[75] However, as firms readjust their optimal bundles, they experience a small annual welfare loss in 2018-2019 for the egress fee ban and 2018-2020 for the increased server interoperability as they readjust their optimal bundles. However, the long run welfare gains more than offset the short run losses. Thus, although the mechanisms differ, my findings are consistent with prior findings in Gans et al. (2023) that an egress fee ban harms consumers in the short run; in the long run, however, lower prices and egress costs raise downstream profits, highlighting the importance of considering changes in price competition as well as industry dynamics when modeling the effects of these policies. In addition, similar to prior literature on IT adoption, I find a J-curve in donwstream firm profits (cf. McElheran et al. (2025)). Though I find immediate increases in adoption and thus median Hicks neutral productivity, those gains do not translate into increases in profits until a few years after the policy implementation (Brynjolfsson et al. (2021)).

[75]In Appendix E, I plot the changes in consumer surplus under the entrant and merger scenarios.



(a) Counterfactual price changes

(b) Counterfactual adoption changes

Figure 9: Counterfactual dynamics: prices and adoption rates

Panel (a) plots the average annual change in price among all suppliers compared to the baseline, including the entrant in the relevant scenario. Panel (b) plots the average annual change in the fraction of firms using cloud computing compared to the baseline. Note that the y-axis is on a log scale. Annual values in (a) and (b) are the mean across 80 block bootstrap samples.

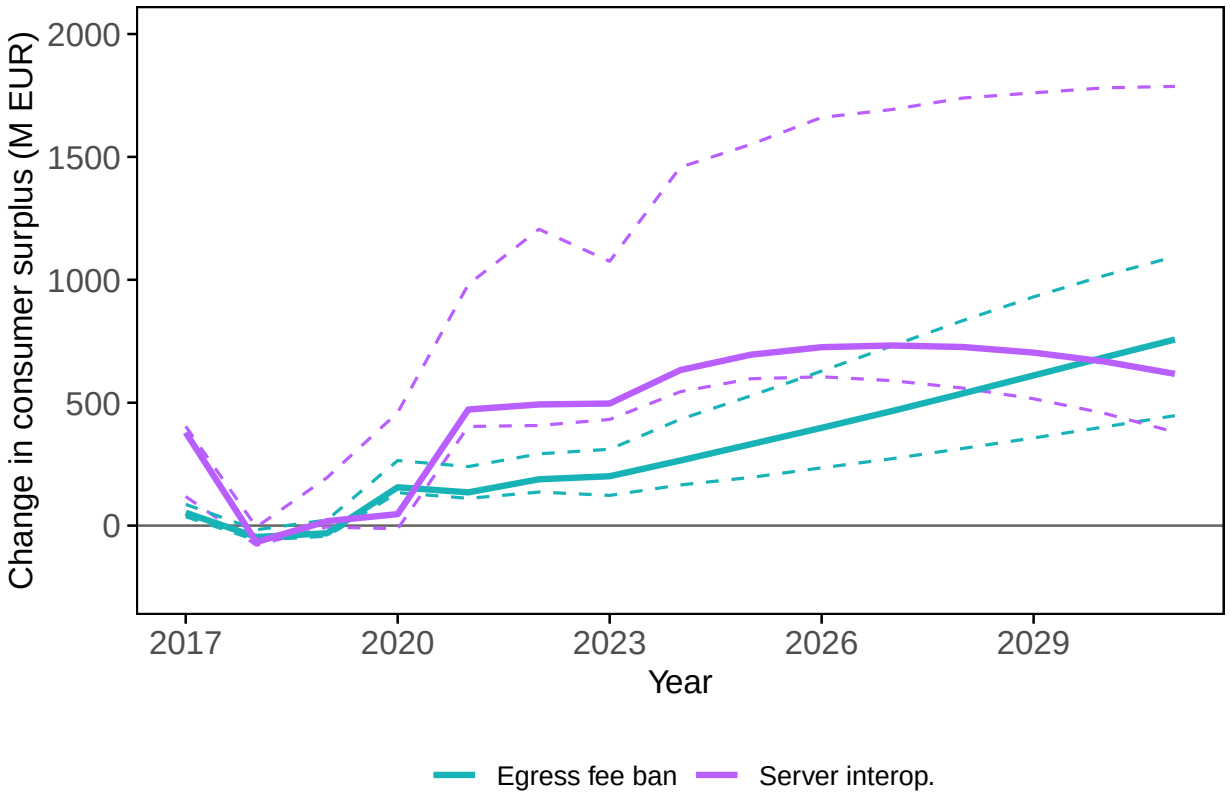


Figure 10: Annual change in consumer surplus under competition policies

Annual change in consumer surplus, 2017-2031, with 95% CI from 80 block bootstrap samples. Values are in real 2015 euros.

8 Conclusion

In this paper, I examine the impact of supplier pricing and market power on downstream firms' input decisions, productivity evolution, and growth in the cloud computing services market. I find that adoption is rising rapidly, and multihoming is common. Despite rapid evolution of the industry, I find that firms' decisions are highly persistent, and supplier concentration is high. To rationalize these findings and explore the impact of recent competition policies, I develop a model of industry dynamics with endogenous IT adoption choices that increase productivity but incur entry, switching, and maintenance costs. Estimating the model, I find gains from adopting cloud computing of between 0.2% and 1.5%. Firms' supplier preferences are driven by complementarities with on-premise technologies as well as price; switching costs are moderate. Using this model, I simulate the impact of a data egress fee ban from the EU Data Act, finding that it increases welfare by €192.6 million per year. I also simulate the impact of better integration between Microsoft server products and non-Microsoft cloud services, finding total welfare gains of €517.2 million per year. Finally, I show that market power matters in this sector - a merger between Microsoft and Google raises virtual machine prices by 25.90% on average, while reducing welfare by €905.3 million per year.

The characteristics in the cloud services market - widespread adoption, high supplier concentration, and multihoming - are common to many software and IT services industries. The modeling approach and insights in the present work may be adapted to study dynamics in other such industries. In particular, in the past 5 years, the market for generative artificial intelligence has been similarly transformed, from firms training their own models to them largely leasing from a few labs. While preliminary evidence suggests that switching barriers are lower than for cloud computing (Demirer et al. (2025)), other features–high barriers to entry among suppliers, multihoming, and rapid adoption–mean that my modeling techniques may be adopted to study this industry.

Before concluding, I briefly discuss extensions and paths for future research. While the present work studies the medium-run impact of pro-competition policies for cloud computing, it focuses on price and cost changes only. In an industry with rapid product improvement, policies that alter incentives to lock in consumers or mandate a certain interoperability level may affect innovation incentives. A study of the effect of such policies on supplier innovation is left for future research. In addition, the present study abstracts away from the impacts of technology adoption on firm entry, exit, and product market competition. Finally, most of these cloud computing firms engage in joint production, supplying not only cloud services but a variety of other computer hardware and software products. The study of joint production

on cloud computing price and feature offerings is left to future research.

References

Acemoglu, D. and Azar, P. D. (2020). Endogenous production networks. *Econometrica*, 88(1):33–82.

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Aguirregabiria, V. and Suzuki, J. (2014). Identification and counterfactuals in dynamic models of market entry and exit. *Quantitative Marketing and Economics*, 12(3):267–304.

Amazon Web Services (2025). Price reduction — aws news blog. Web page, blog category. Accessed: 2025-09-25.

Arcidiacono, P. and Miller, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867.

Asker, J., Collard-Wexler, A., and De Loecker, J. (2014). Dynamic inputs and resource (mis) allocation. *Journal of Political Economy*, 122(5):1013–1063.

Atalay, E., Hortaçsu, A., and Syverson, C. (2014). Vertical integration and input flows. *American Economic Review*, 104(4):1120–1148.

Autorité de la concurrence (2023). Opinion 23-a-08 of 29 june 2023 on competition in the cloud sector. Opinion 23-A-08, Autorité de la concurrence. English version published 28 September 2023.

Aw, B. Y., Roberts, M. J., and Xu, D. Y. (2011). R&d investment, exporting, and productivity dynamics. *American Economic Review*, 101(4):1312–1344.

Azar, P. D. (2021). Moore's law and economic growth. *FRB of New York Staff Report*, (970).

Baqaee, D. R. and Farhi, E. (2018). Macroeconomics with heterogeneous agents and input-output networks. Technical report, National Bureau of Economic Research.

Barefoot, K., Curtis, D., Jolliff, W., Nicholson, J. R., Omohundro, R., et al. (2018). Defining and measuring the digital economy. *US Department of Commerce Bureau of Economic Analysis, Washington, DC*, 15:210.

Bertolotti, F., Lanteri, A., and Villa, A. (2025). Investment-goods market power and capital accumulation.

Besiroglu, T., Emery-Xu, N., and Thompson, N. (2024). Economic impacts of ai-augmented r&d. *Research Policy*, 53(7):105037.

Bloom, N. (2009). The impact of uncertainty shocks. *econometrica*, 77(3):623–685.

Bloom, N., Draca, M., and Van Reenen, J. (2016). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *The review of economic studies*, 83(1):87–117.

Bloom, N. and Pierri, N. (2018). Cloud computing is helping smaller, newer firms compete. *Harvard Business Review*, 94(4).

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

Boudreau, B. (2022). Price erosion remains the universal norm.

Brand, J. M., Demirer, M., Finucane, C., and Kreps, A. A. (2024). Firm productivity and learning in the digital economy: Evidence from cloud computing. Technical report, National Bureau of Economic Research.

Bresnahan, T. F. (1982). The oligopoly solution concept is identified. *Economics letters*, 10(1-2):87–92.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of econometrics*, 65(1):83–108.

Brynjolfsson, E. and Hitt, L. (1996). Paradox lost? firm-level evidence on the returns to information systems spending. *Management science*, 42(4):541–558.

Brynjolfsson, E. and Hitt, L. M. (2003). Computing productivity: Firm-level evidence. *Review of economics and statistics*, 85(4):793–808.

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.

Byrne, D., Corrado, C., and Sichel, D. E. (2018). The rise of cloud computing: minding your p's, q's and k's. Technical report, National Bureau of Economic Research.

Chari, V. V. and Hopenhayn, H. (1991). Vintage human capital, growth, and the diffusion of new technology. *Journal of political Economy*, 99(6):1142–1165.

Collard-Wexler, A. and De Loecker, J. (2015). Reallocation and technology: Evidence from the us steel industry. *American Economic Review*, 105(1):131–171.

Competition and Markets Authority (2024). Egress fees working paper. Working paper, cloud services market investigation, Competition and Markets Authority. Consultative working paper published 23 May 2024.

Competition and Markets Authority (2025a). Appendix a: Demand for cloud services and how customers purchase cloud services. Appendix to final report, cloud services market investigation, UK Competition and Markets Authority. Appendix A to the CMA's final decision on the market investigation into public cloud infrastructure services.

Competition and Markets Authority (2025b). Appendix i: Prevalence of multi-cloud. Appendix to final decision report, cloud services market investigation, UK Competition and Markets Authority. Appendix I — Analysis of providers' and customers' data on how common multi-cloud usage is.

Competition and Markets Authority (2025c). Appendix m: Egress fees – analysis of customers' cost scenarios. Appendix to final decision report, cloud services market investigation, UK Competition and Markets Authority. Appendix M — Egress fees: analysis of customers' cost scenarios.

Competition and Markets Authority (2025d). Appendix o: Customer views on egress fees. Appendix to final decision report, cloud services market investigation, UK Competition and Markets Authority. Appendix O — Customer views on egress fees.

Competition and Markets Authority (2025e). Appendix q: Egress fees — analysis of cloud providers' cost of egress. Appendix to final decision report, cloud services market investigation, UK Competition and Markets Authority. Appendix Q — Egress fees: analysis of cloud providers' cost of egress.

Competition and Markets Authority (2025f). Appendix t: Licensing analysis. Appendix to final decision report, cloud services market investigation, UK Competition and Markets Authority. Appendix T to the final decision report on the Cloud Services Market Investigation.

Competition and Markets Authority (2025g). Appendix w: Remedies not taken forward in this market investigation. Appendix to final decision report, cloud services market investigation, UK Competition and Markets Authority. Remedies Appendix — potential remedies considered but not pursued under CMA powers.

Competition and Markets Authority (2025h). Chapter 3: Competitive landscape. In *Cloud services market investigation: Final decision report*, report chapter 3. Final Decision Report of the CMA.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633.

Cortez, E., Bonde, A., Muzio, A., Russinovich, M., Fontoura, M., and Bianchini, R. (2017). Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 153–167.

Corts, K. S. (1999). Conduct parameters and the measurement of market power. *Journal of Econometrics*, 88(2):227–250.

Crawford, G. S. and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, 102(2):643–685.

Datacenters.com (2025). Search data centers: Colocation, bare metal, cloud. Online database / marketplace. Accessed: 2025-11-05.

de Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics*, 5(3):1–21.

De Loecker, J., Goldberg, P. K., Khandelwal, A. K., and Pavcnik, N. (2016). Prices, markups, and trade reform. *Econometrica*, 84(2):445–510.

De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review*, 114(1):199–251.

Demirer, M., Fradkin, A., Tadelis, N., and Peng, S. (2025). The emerging market for intelligence: Pricing, supply, and demand for llms. Technical report, National Bureau of Economic Research.

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.

DeStefano, T., Kneller, R., and Timmis, J. (2023). Cloud computing and firm growth. *Review of Economics and Statistics*, pages 1–47.

Do, J. and Miklós-Thal, J. (2025). Price discrimination against multi-clouders. *Available at SSRN*.

Dos Santos, M. and Ovsyannikova, T. (2023). Solutions saas : quelle comptabilisation dans les comptes établis en normes françaises? Web article. Accessed: 2025-09-19.

Dubé, J.-P., Hitsch, G. J., and Rossi, P. E. (2009). Do switching costs make markets less competitive? *Journal of Marketing research*, 46(4):435–445.

EC2Instances.info Contributors (2025). Ec2instances.info: Easy amazon ec2 instance comparison. Accessed: 2025-09-15.

Emeras, J., Varrette, S., Plugaru, V., and Bouvry, P. (2016). Amazon elastic compute cloud (ec2) versus in-house hpc platform: A cost analysis. *IEEE Transactions on Cloud Computing*, 7(2):456–468.

Ericson, R. and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of economic studies*, 62(1):53–82.

Eurostat / European Commission (2023). 45 News release, Eurostat. Accessed: 2025-09-25.

Farrell, J. and Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. *Handbook of industrial organization*, 3:1967–2072.

Flexera (2019). Rightscale 2019 state of the cloud report from flexera. Technical report, Flexera. Survey of 786 technical professionals; licensed under Creative Commons Attribution4.0 International License.

Flexera (2020). Rightscale 2020 state of the cloud report from flexera. Technical report, Flexera.

Foster, L., Haltiwanger, J., and Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425.

Galindo-Rueda, F. and Verger, F. (2016). Oecd taxonomy of economic activities based on r&d intensity. Technical Report 2016/04, OECD Publishing, Paris.

Gans, J., Hervé, M., and Masri, M. (2023). Economic analysis of proposed regulations of cloud services in europe. *European Competition Journal*, 19(3):522–568.

Garcia-Osipenko, M., Vreugdenhil, N., and Zahur, N. B. (2025). Endogenous rigidities and capital misallocation: Evidence from containerships. *Available at SSRN 5217587*.

Gavazza, A. (2011). Leasing and secondary markets: Theory and evidence from commercial aircraft. *Journal of Political Economy*, 119(2):325–377.

Goettler, R. L. and Gordon, B. R. (2011). Does amd spur intel to innovate more? *Journal of Political Economy*, 119(6):1141–1200.

Gortmaker, J. (2024). Open source software policy in industry equilibrium. Technical report, Working paper, Tech. Rep.

Gowrisankaran, G. and Rysman, M. (2012). Dynamics of consumer demand for new durable goods. *Journal of political Economy*, 120(6):1173–1219.

Gowrisankaran, G. and Schmidt-Dengler, P. (2024). A computable dynamic oligopoly model of capacity investment. Technical report, National Bureau of Economic Research.

HardDriveBenchmark.net (n.d.). Hard drive benchmark. Online database. Accessed: 2025-09-19.

Heckman, J. J. et al. (1981). Statistical models for discrete panel data. *Structural analysis of discrete data with econometric applications*, 114:178.

Honka, E. (2014). Quantifying search and switching costs in the us auto insurance industry. *The RAND Journal of Economics*, 45(4):847–884.

Hornbeck, R., Hsu, S. H.-M., Humlum, A., and Rotemberg, M. (2024). Gaining steam: Incumbent lock-in and entrant leapfrogging. Technical report, National Bureau of Economic Research.

Humlum, A. (2022). *Robot adoption and labor market dynamics*. Rockwool Foundation Research Unit Berlin, Germany.

Ifrach, B. and Weintraub, G. Y. (2017). A framework for dynamic oligopoly in concentrated industries. *The Review of Economic Studies*, 84(3):1106–1150.

Igami, M. and Uetake, K. (2020). Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016. *The Review of Economic Studies*, 87(6):2672–2702.

Japan Fair Trade Commission (2022). Report on trade practices in cloud services sector. Report on fact-finding survey on trade practices by digital platform operators, Japan Fair Trade Commission. Report on Fact-Finding Survey on Trade Practices by Digital Platform Operators.

Jin, W. and Bai, J. J. (2022). Cloud adoption and firm performance: Evidence from labor demand. *Available at SSRN 4082436*.

Jin, W. and McElheran, K. (2024). Economies before scale: It strategy and performance dynamics of young us businesses. *Management Science*.

Kearns, A. J. (2024). Broadband deployment in equilibrium: Is competition the solution to the digital divide? Working paper, Federal Trade Commission.

Khajeh-Hosseini, A., Greenwood, D., and Sommerville, I. (2010). Cloud migration: A case study of migrating an enterprise it system to iaas. In *2010 IEEE 3rd International Conference on cloud computing*, pages 450–457. IEEE.

Kilcioglu, C. and Rao, J. M. (2016). Competition on price and quality in cloud computing. In *Proceedings of the 25th international conference on world wide web*, pages 1123–1132.

Korea Fair Trade Commission (2022). Kftc announces results of cloud service market study. `https://www.ftc.go.kr/eng/downloadBbsFile.do?atchmnflNo=17414`. Press release, December 28, 2022 (hereinafter "S. Kor. FTC Cloud Study").

Korinek, A. and Vipra, J. (2025). Concentrating intelligence: scaling and market structure in artificial intelligence. *Economic Policy*, 40(121):225–256.

Lashkari, D., Bauer, A., and Boussard, J. (2024). Information technology and returns to scale. *American Economic Review*, 114(6):1769–1815.

Mairesse, J., Mohnen, P., Kremp, E., and KREMP, E. (2005). The importance of r&d and innovation for productivity: A reexamination in light of the french innovation survey. *Annales d'Economie et de Statistique*, pages 487–527.

McCallum, J. C. and of Labor Statistics, U. B. (2024). Disk. Based on: John C. McCallum (2023), *Price and Performance Changes of Computer Technology with Time*; U.S. Bureau of Labor Statistics (2024), *US Consumer Prices*; with minor processing by Our World in Data.

McElheran, K., Yang, M.-J., Kroff, Z., and Brynjolfsson, E. (2025). *The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve (s)*. US Census Bureau, Center for Economic Studies Suitland, MD, USA.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society*, pages 995–1026.

MultiCloud.Tips (n.d.). Case studies of successful multi-cloud deployments. Web Article. Accessed: 2025-09-16.

North, M. A., Cavallaro, R., and Davison, D. (2019). The total economic impact™ of microsoft azure iaas: Cost savings and business benefits enabled by azure infrastructure-as-a-service (iaas). Technical report, Forrester Consulting, commissioned by Microsoft. Forrester Research, Inc.

Ofcom (2023). Cloud services market study: Interim report. Technical report, Office of Communications (Ofcom). Accessed: 25 February 2026.

Pakes, A. and Pollard, D. (1989). Simulation and the asymptotics of optimization estimators. *Econometrica: Journal of the Econometric Society*, pages 1027–1057.

Pan Fang, T. and Greenstein, S. (2025). Where the cloud rests: The economic geography of data centers. *Strategy Science*.

Park, M. (2011). The economic impact of wireless number portability. *The Journal of Industrial Economics*, 59(4):714–745.

PassMark Software (2026). Cpu benchmarks. `https://www.cpubenchmark.net/`. Accessed: 2026-04-17.

Peters, B., Roberts, M. J., Vuong, V. A., and Fryges, H. (2017). Estimating dynamic r&d choice: an analysis of costs and long-run benefits. *The RAND Journal of Economics*, 48(2):409–437.

Phillips, G. and Owens, S. (2022). The total economic impact™ of aws cloud operations: Cost savings and business benefits enabled by operating on aws. Technical report, Forrester Consulting, commissioned by AWS. Forrester Research, Inc.

Piper, S. and Cohen, A. (2023). The state of the cloud 2023: The top cloud security threats to be aware of in 2023. Blog post, Wiz. Accessed: 2025-09-18.

Polyakova, M. (2016). Regulation of insurance with adverse selection and switching costs: Evidence from medicare part d. *American Economic Journal: Applied Economics*, 8(3):165–195.

Reiss, P. C. and Wolak, F. A. (2007). Structural econometric modeling: Rationales and examples from industrial organization. *Handbook of econometrics*, 6:4277–4415.

Rubens, M. (2025). Labor-market power, deadweight loss, and technology adoption.

Rubens, M., Wu, Y., and Xu, M. (2024). Exploiting or augmenting labor. Technical report, Working Paper.

Seagate Technology (2020). Rethink data: Put more of your business data to work—from edge to cloud. Industry report, Seagate Technology. With research and analysis by IDC; based on a global survey of 1,500 enterprise respondents.

Shcherbakov, O. (2016). Measuring consumer switching costs in the television industry. *The RAND Journal of Economics*, 47(2):366–393.

Statista (2025). Public cloud – infrastructure as a service (iaas), france. Online market forecast. Accessed: 2025-09-19.

Tambe, P. and Hitt, L. M. (2012). The productivity of information technology investments: New evidence from it labor data. *Information systems research*, 23(3-part-1):599–617.

Thompson, N. (2017). The economic impact of moore's law: Evidence from when it faltered. *Available at SSRN 2899115*.

TRENDFOCUS, Inc. (2019). Cloud, hyperscale, & enterprise storage service: Cq1 '19 mid-quarter update. Industry report, TRENDFOCUS, Incorporated, Cupertino, CA. Published April 4, 2019. Authors: Mark Geenen, John Kim, John Chen, Don Jeanette, Jason Lervold.

U.S. Department of Justice (2001). United states v. microsoft corporation: Competitive impact statement. Competitive impact statement, U.S. Department of Justice.

Vantage (2022). Cloud cost report: Q3 2022. Industry report, Vantage. Analysis of anonymized cloud spending data from thousands of AWS accounts.

Viard, V. B. (2007). Do switching costs make markets more or less competitive? the case of 800-number portability. *The RAND Journal of Economics*, 38(1):146–163.

Wang, S., LaRiviere, J., and Kannan, A. (2019). Spatial competition and missing data: an application to cloud computing. *Work. Pap*, pages 1–23.

Yang, C. (2020). Vertical structure and innovation: A study of the soc and smartphone industries. *The RAND Journal of Economics*, 51(3):739–785.

Yeo, J. and Miller, D. P. (2018). Estimating switching costs with market share data: an application to medicare part d. *International Journal of Industrial Organization*, 61:459–501.

Zhuo, R. and Tiew, A. (2025). Industrial policy, location choice, and firm performance in high-tech manufacturing. *Location Choice, and Firm Performance in High-Tech Manufacturing (August 01, 2025)*.

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., and Dinlersoz, E. (2021). Advanced technologies adoption and use by us firms: Evidence from the annual business survey. Technical report, National Bureau of Economic Research.

Appendix

A Additional dataset details **61**
A.1 Dataset description . . . 61
A.2 Matching procedure . . . 62
A.3 Cloud products observed in data . . . 63
A.4 Additional summary statistics . . . 64
A.5 Capital stock computation . . . 64
A.6 Computing VM prices and usage . . . 65
A.6.1 VM usage dataset description . . . 65
A.6.2 Usage and price merge . . . 68
A.6.3 Computing virtual machine benchmarks . . . 69
A.6.4 Datacenter locations . . . 72
A.6.5 Feature and region merge . . . 74
A.7 Microsoft server products . . . 74
A.8 Calibrating marginal costs . . . 76
A.8.1 vCPU marginal cost . . . 76
A.8.2 Egress marginal cost . . . 77
A.9 Evidence of worldwide IaaS market concentration . . . 77
A.10 Difference-in-differences regression on cloud spot prices . . . 78

B Modeling details appendix **79**
B.1 Capital investment details . . . 80
B.2 Bundle choice details . . . 81
B.3 Supplier pricing details . . . 82
B.4 A model of capital and data . . . 84

C Estimation appendix **85**
C.1 Exogenous state transitions . . . 85
C.2 Dynamic parameter estimation details . . . 86
C.3 Computational details . . . 90

D Results appendix **91**
D.1 Additional production function estimates . . . 91
D.1.1 Differential effects of single- and multihoming . . . 92
D.1.2 Translog production function . . . 93

D.1.3 Learning-by-doing and productivity 94
D.1.4 Total investment proxy variable 95
D.1.5 Including exporting and market share quality shifters 96
D.2 Steady state productivity distribution for multihoming firms 98
D.3 Cost distribution for remaining sectors 99
D.4 Comparison to industry estimates of IaaS costs and benefits 100

E Additional counterfactual results **101**
E.1 Stylized model of cloud adoption and pricing 101
E.1.1 Effect of egress fees 102
E.1.2 Effect of improved server integration 103
E.2 Deriving the counterfactual invest-harvest parameter 104
E.3 Additional counterfactual results 105
E.4 Additional counterfactual figures 106

F Further institutional details **109**
F.1 The Infrastructure-as-a-Service market 109
F.2 Competition policy for cloud computing services 112

A Additional dataset details

This appendix provides additional details on my dataset construction. Appendix A.1 provides an overview of the datasets. Appendix A.2 describes the matching procedure I use to combine the Orbis and Aberdeen data. Appendix A.3 lists the cloud computing products observed in my sample. Appendix A.4 provides additional summary statistics for the full sample of Orbis firms. Appendix A.5 describes how I compute capital stocks and investment rates. Appendix A.6 provides details on the virtual machine price and quality computation. Appendix A.7 lists the Microsoft server-related products in the data. Appendix A.8 describes marginal cost calibration. Appendix A.9 provides evidence of market concentration among IaaS suppliers worldwide. Finally, Appendix A.10 provides evidence against lowering egress fees mattering for switching in the short run, leveraging Google's partial elimination of egress fees in January 2024.

A.1 Dataset description

Downstream firm data. My primary firm data comes from Bureau van Dijk's Orbis dataset, which covers unconsolidated firm financial data for a large fraction of public and private French firms. Prior research has found a high level of coverage of the dataset,[76] and I find that the dataset covers and 71.1% and 68.5% of output in the manufacturing and services sectors, respectively. Once I merge the Aberdeen data to fill in missing revenue, which comes from Dun & Bradstreet, coverage rates improve to 76.2% in manufacturing and 89.3% in services.

The data on cloud and other IT investments comes from Aberdeen (formerly Harte Hanks), a market research firm, which collects IT data to sell to computer software and hardware vendors. Aberdeen's data collection relies on a series of over 10,000 monthly surveys, from which they obtain data on firm IT stock and software vendor choice at the establishment level (Aberdeen Group (2017)). In addition, using these survey responses, they estimate the value of IT investments at each establishment, and rely on web scraping and interpolation to impute missing values, which can introduce measurment error (Brynjolfsson and Hitt (2003); Demirer et al. (2024)). Historically, the firm has targeted for survey all firms with 100 or more employees, though recently it has expanded to include smaller firms. From 1998-2007, Aberdeen had a 45.1% response rate among European firms, increasing in firm size (Bloom et al. (2012)). By comparison, my sample has a 24.2% match rate among *all* French firms in Orbis.

[76] Asker et al. (2014)

Cloud computing price and feature data. My cloud computing price data comes from the list price pages for each cloud provider, archived on the Wayback Machine. Since Amazon's prices are not preserved because dynamic loading via AJAX is not archived, I use archived versions of `ec2instances.info` (EC2Instances.info Contributors (2025)). I gather data on prices and vCPU count at the quarterly level, using data as close to the end of each quarter as possible and take the average to get annual values. Next, I gather CPU price and benchmark performance data from Passmark, which provides a common unidimensional rating system for CPUs. 4To get datacenter locations, I use location data from `Datacenters.com`. As described below, to compute average annual provider price and feature data, I first merge the set of VMs available at each data center year to representative monthly usage data from Microsoft Azure (Cortez et al. (2017); Microsoft Azure (2019)). As described below, I exact match Microsoft virtual machines and match Google and Amazon virtual machines to account for slight variation in the specs offered for similar products. Reassuringly, I find the variance in the average price per vCPU hour across customers to be quite small across both usage datasets, allowing me to take weighted averages of prices and features at the datacenter-year level. Finally, I assume that each firm uses the datacenter closest to its headquarters to minimize latency.

Other sources. Price deflators largely come from Eurostat, with the exception of the price index for computer storage, which comes from the US BLS.

A.2 Matching procedure

As the firmographic datasets do not have any common firm identifiers, I use other identifying information to merge the datasets. My matching algorithm is as follows. First, to ensure that I am matching at the same level of ownership, I restrict my Orbis sample to firms with unconsolidated financial data.[77] For Aberdeen, I merge all establishments that have the same enterprise ID and name.[78] Second, following the approach in Hazell et al. (2022), I clean firm names by converting them to lowercase, removing common words and abbreviations like "company" and "ltd." I then run the following matching algorithm, modified from Demirer et al. (2024):

1. Match unique exact firm names
2. Match firms with the same URL
3. Match firms with the same phone number

[77] CONSCODE of "U1" or "U2."
[78] This implicitly ignores ownership structure.

4. Match firms in the same postal code, and the Jaro-Winkler name distance is less than 0.1

5. Match firms in the same city, and the name distance is less than 0.08

6. Match firms in the same province, and the name distance is less than 0.07

7. Match firms in the same country, and the name distance is less than 0.065

For city and province matches, since a subsidiary firm may have multiple headquarters listed, I am unable to determine which headquarters location is the corresponding one to that in Orbis. Therefore, I consider a match to be positive if the match conditions are met and any establishment is in the same city or province as in the Orbis data.[79] Using this procedure, I am able to match 69.5% of firms in the Aberdeen sample and 24.2% of firms in the Orbis sample. Manual spot checks ensure the accuracy of the merged dataset.

A.3 Cloud products observed in data

Table A.1 lists the products observed in Aberdeen's surveys. Note that the level of detail varies by cloud provider, with only Amazon and Microsoft listing both storage and computation, restricting the level of detail at which I am able to estimate egress fees. However, egress fees apply to both virtual machine storage and dedicated storage blobs, so my assumption that all firms exiting a specific IaaS provider are subject to egress fees seems reasonable.

Table A.1: Cloud products grouped by manufacturer

Cloud provider	**Products**
Amazon	Amazon S3, Amazon EC2, Amazon Web Services
Microsoft	Azure, Azure Blob Storage, Azure Virtual Machines
Google	Cloud Platform, Compute Engine

List of all product names that have appeared it at least one year of Aberdeen's data from 2017-2021. Names have been standardized across survey waves.

[79] This may introduce bias if Aberdeen does not sample every establishment for a given firm. Until 2009, Aberdeen included data on total number of establishments at each firm as well as the fraction surveyed. For each year in the US sample until 2009, the mean firm had data on over 95% of establishments, and the median firm had data on 100% of establishments.

A.4 Additional summary statistics

I present summary statistics for the full Orbis sample in the following table.[80] Confirming prior research (e.g., Bloom et al. (2012)) the Aberdeen sample is weighted toward larger firms, with average revenue that is 3.28 times the overall sample and average employee count that is 1.44 times the overall sample. Further, firms in the matched sample are around 10 years older than firms in the full sample.

Table A.2: Full Orbis sample summary statistics

Variable	Observations	Mean	25th Percentile	Median	75th Percentile
Revenue	2,962,121	7.47	0.08	0.39	1.71
Labor	1,285,080	48.59	3.00	9.00	20.00
Capital	3,375,219	5.37	0.01	0.10	0.44
Materials	2,707,271	6.07	0.04	0.18	0.93
Firm Age	3,375,358	14.52	4.00	10.00	21.00

Notes: The units for revenue, capital, and materials are in millions of real 2015 euros. The units for labor are in employees. The unit for firm age is years.

A.5 Capital stock computation

Because I do not have enough observations per firm in Orbis, I use the book value of the capital stock for each firm. In this appendix, I describe how I recover total investment and non-IT investment. First, I deflate all variables using capital-type specific deflators from Eurostat. Next, Orbis contains depreciation data for a subset of firms. For the remaining firms, I compute

$$\delta_{\iota t} = \mathbb{E}_{\iota t}\left[1 + \frac{I_{ft}}{K_{ft}} - \frac{p_{\iota t+1}K_{ft+1}}{p_{\iota t}K_{ft}}\right]$$

where $\delta_{\iota t}$ is then the average annual depreciation rate for industry ι at the 4-digit NACE code level. Using this, I compute the value of the remaining investment stock. Finally, Aberdeen has values for establishment level IT capital stocks (storage, hardware, and software), which I sum up to the firm level, deflate by their relevant Eurostat deflators, and subtract from I_{ft} for each f, t.

The exception is the stock of data. Aberdeen only measures on-premise storage investment. Thus, I limit my calculations to the pre-cloud period (2015-2016). I make the assumption (cf. Lashkari et al. (2024)) that, in 2025, firms' data stocks are given by

[80] Revenue, labor, and industry code data are sourced from Aberdeen if the reelevant data is missing in Orbis. I filter the data to manufacturing and services excluding financial services and healthcare.

$$D_{it} = \frac{I^D_{ft}}{I_{ft}} \cdot K_{ft}$$

where I^D_{ft} is the investment in data, given by Aberdeen in nominal euros and converted to petabytes (PB) using the BLS computer storage deflator (McCallum and of Labor Statistics (2024)). For 2016, then, the data stock is

$$D_{ft} = (1 - \delta^D)D_{ft-1} + I^D_{ft}$$

where $\delta^D = 0.36$ from Bloom et al. (2012).

A.6 Computing VM prices and usage

This section describes my cloud computing product dataset in more detail, including the process to merge data sources and compute datacenter-year level prices, costs, and features.

A.6.1 VM usage dataset description

To compute the average price paid at the datacenter-year level, I utilize anonymized virtual machine usage data from Microsoft Azure (Microsoft Azure (2019); Cortez et al. (2017)). The first trace data is from 2017 and contains 30 consecutive days of VM usage for 5,958 subscriptions and 2,013,767 VMs. The second is from 2019 and contains 30 consecutive days of VM usage for 6,689 subscriptions and 2,695,548 VMs. This dataset has been used in prior economics research to understand firm learning-by-doing in virtual machine utilization (Brand et al. (2024)). The data contains the category (interactive vs delay-insensitive), categories for number of cores (0-2, 4-8, etc.), and memory size (0-2 GB, 2-4 GB, etc.).[81] To merge the data, for each region, I assume that firms pay the average price among all VMs in a given category.

Below, I present summary statistics for VM prices. Let w_{fjt} be the usage (measured in seconds) of VM instance type v used by firm f in year t, let p_{fvt} be the corresponding price, and let $W_{ft} = \sum_v w_{fvt}$. Then I can decompose the total (usage-weighted) variance of VM prices into the variance within firms and the variance between firms as follows:

[81] The public data is from either internal Microsoft-owned firms' IaaS usage or external firms' usage of Microsoft PaaS. The sample has been selected to be representative of external Azure users in the month-region of study. For further details on the 2017 trace, see Cortez et al. (2017), and on the 2019 trace, see Microsoft Azure (2019).

$$\sigma^2_{t,\text{total}} = \underbrace{\frac{\sum_f W_{ft} \frac{\sum_j w_{fvt}\left(p_{fvt}-\bar{p}_{ft}\right)^2}{\sum_j w_{fvt}}}{\sum_f W_{ft}}}_{\sigma^2_{t,\text{within}}} + \underbrace{\frac{\sum_f W_{ft}\left(\bar{p}_{ft}-\bar{p}_t\right)^2}{\sum_f W_{ft}}}_{\sigma^2_{t,\text{between}}}$$

$$\text{where} \quad \bar{p}_{ft} = \frac{\sum_v w_{fvt}\, p_{fvt}}{\sum_v w_{fvt}}$$

$$\bar{p}_t = \frac{\sum_f \sum_v w_{fvt}\, p_{fvt}}{\sum_f \sum_v w_{fvt}}.$$

Results are presented for the West Europe region in Table A.3. While variance between firms is similar in both time periods, the variance within firms increased, likely due to an expanding set of VM offerings in 2019 (139) compared to 2017 (65). In addition, Cortez et al. (2017) and Microsoft Azure (2019) show an increasing rate of per-VM CPU utilization[82], indicating firms in 2019 are better than firms in 2017 at matching VM compute to appropriate workloads (cf. Brand et al. (2024)).

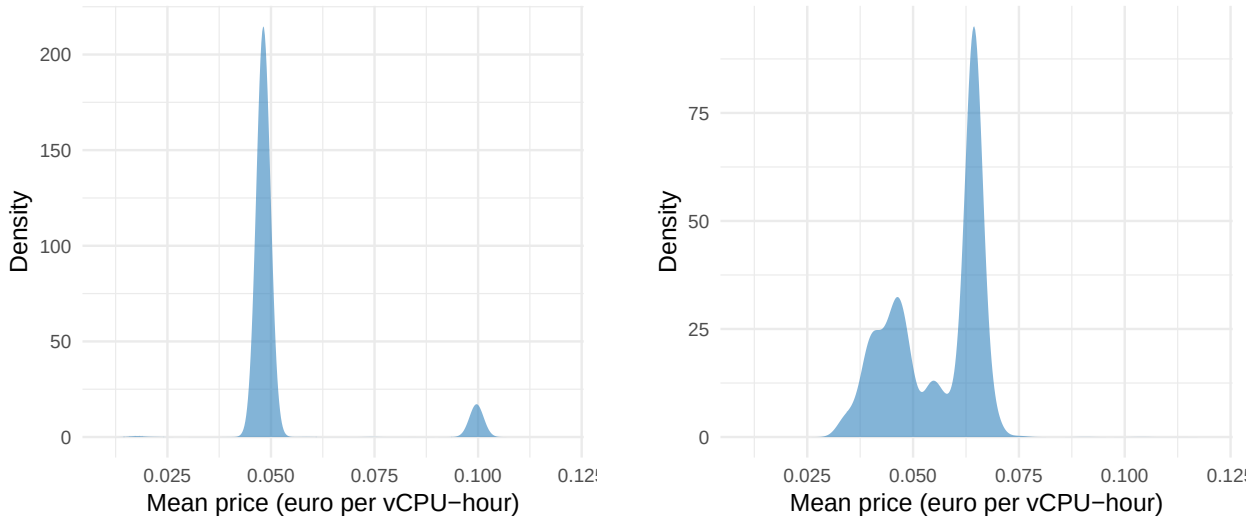
Next, Figures A.1 and A.2 display the density of vCPU prices between and within firms, respectively. While there is a narrow range of prices that most firms pay in both years, there are several smaller modes, especially in 2019. Likewise, the within-firm distribution is either degenerate for smaller firms or approximately normally distributed. The mean price paid does not vary much by firm size.

Table A.3: Variance decomposition of VM prices per core-hour

Variance	2017	2019
Between variance ($\sigma^2_{between}$)	1.84×10^{-4}	1.06×10^{-4}
Within variance (σ^2_{within})	2.30×10^{-6}	4.55×10^{-5}
Total variance ($\sigma^2_{between} + \sigma^2_{within}$)	1.86×10^{-4}	1.67×10^{-4}

Notes: Weighted variance of price per vCPU-hour within and between firms for an anonymous Microsoft Azure region in the period November 16, 2016 to February 16, 2017 and some period ending before July 2019. For the former, price data from February 2017 is used. For the latter, price data from March 2019 is used. In 2019, 1.7% of firm-VM type observations are dropped due to the lack of a corresponding VM instance type in West Europe during that time.

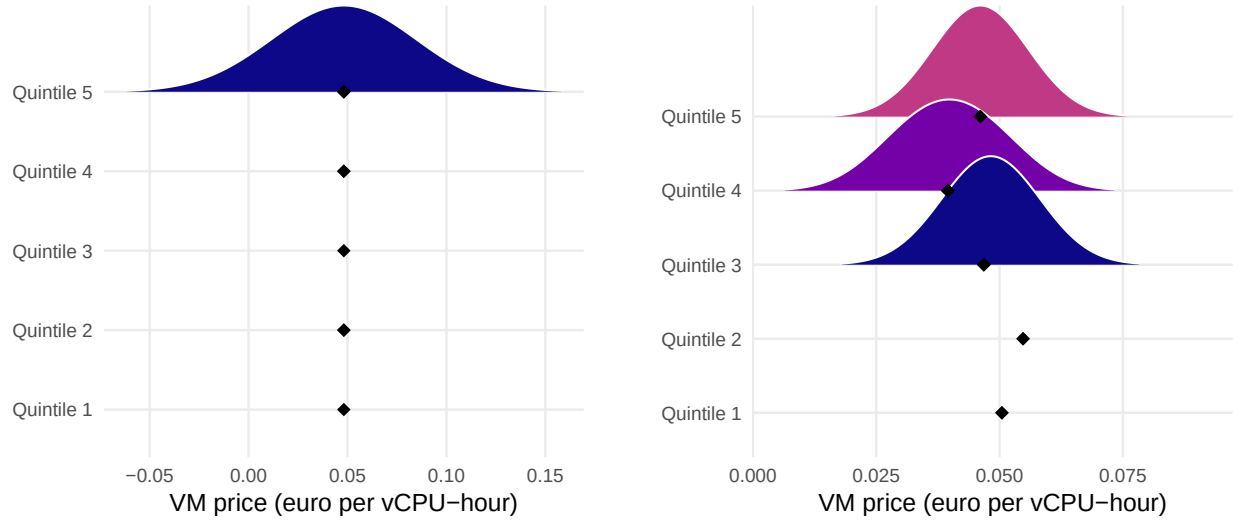
[82]That is, the fraction of CPU resources used by user workloads over the lifetime of the VM.



(a) 2017 (b) 2019

Figure A.1: Density of vCPU prices between firms

Note: Weighted average of firm level vCPU prices in euros.



(a) 2017 (b) 2019

Figure A.2: Density of vCPU prices within firm by quintile

Note: Weighted average of firm level vCPU prices in real euros. Firms are selected as the median firm in each quintile, measured by total VM usage in seconds.

A.6.2 Usage and price merge

Given that I only observe prices for a single Azure region in 2017 and in 2019, I assume that all firms use the same distribution of virtual machines for every supplier j in each region l:

Assumption 2. *Let q_{vlt} denote the quality of instance type v in region l at time t and let $\mathcal{F}_j^w(q_{vlt})$ and $\mathcal{F}_j^b(q_{vlt})$ denote the within and between consumer distribution of q_{vlt}, respectively. Then*

$$\begin{aligned}
\mathcal{F}_j^w(q_{vlt}) &= \mathcal{F}_{j'}^w(q_{vlt}) \\
\mathcal{F}_j^w(q_{vlt}) &= \mathcal{F}_j^w(q_{vl't}) \\
\mathcal{F}_j^b(q_{vlt}) &= \mathcal{F}_{j'}^b(q_{vlt}) \\
\mathcal{F}_j^b(q_{vlt}) &= \mathcal{F}_j^b(q_{vl't}).
\end{aligned}$$

To merge Microsoft Azure price data, I perform an exact match on VMs based on number of cores and memory, taking the mean price per core hour within each category. As the VMs offered by Amazon and Google are not exact matches for Microsoft VMs, I employ a fuzzy matching procedure in which I match exactly on CPU core count and match memory size in GB according to the size bins $(2^k, 2^{k+1}]$ for integers $k \in [0, 7]$.[83] I match the 2017 Azure trace dataset to 2016-2017 pricing data and the 2019 trace dataset to 2018-2021 pricing data.[84]

The following Table A.4 displays, for each supplier, the annual weighted average price per core hour in nominal euros and fraction of VMs in the trace data that matched successfully to price data. For each supplier-year, I display the average, min, and max prices and match rates across locations l. Google's match rate is lowest, likely due to the fact that it offered fewer SKUs during this period than its competitors. This problem is severe in 2017 but ameliorated in the rest of the sample.

[83] For example, Microsoft's basic A1 VM has 1.88 GB of RAM while the equivalent Amazon SKU, the t2.small, has 2.15 GB.

[84] The 2019 trace data has a higher match rate to 2018 Azure pricing data than the 2017 trace data. To ensure consistency, I adopt this assumption across all providers.

Table A.4: Virtual machine price summary statistics

		Price per core hour			Usage match rate		
Provider	Year	Mean	Min	Max	Mean	Min	Max
Amazon	2017	0.0559	0.0492	0.0626	0.7510	0.5030	1.0000
Amazon	2018	0.0499	0.0380	0.0605	0.8040	0.7170	0.9790
Amazon	2019	0.0575	0.0551	0.0602	0.8040	0.7170	0.9790
Amazon	2020	0.0533	0.0438	0.0577	0.8480	0.7170	0.9790
Amazon	2021	0.0434	0.0365	0.0515	0.9790	0.9790	0.9790
Google	2017	0.0649	0.0608	0.0670	0.0691	0.0691	0.0691
Google	2018	0.0487	0.0449	0.0525	0.8560	0.8560	0.8560
Google	2019	0.0477	0.0379	0.0601	0.8560	0.8560	0.8560
Google	2020	0.0440	0.0394	0.0468	0.8560	0.8560	0.8560
Google	2021	0.0421	0.0368	0.0463	0.8560	0.8560	0.8560
Microsoft	2017	0.0531	0.0531	0.0531	1.0000	1.0000	1.0000
Microsoft	2018	0.0595	0.0552	0.0724	0.9790	0.9790	0.9790
Microsoft	2019	0.0643	0.0553	0.0810	0.9790	0.9790	0.9790
Microsoft	2020	0.0626	0.0476	0.0769	0.9790	0.9790	0.9790
Microsoft	2021	0.0572	0.0407	0.0712	0.9790	0.9790	0.9790

A.6.3 Computing virtual machine benchmarks

This appendix describes how I compute the average VM quality using CPU benchmarks, which provide a unidimensional indicator of a CPU's performance (Goettler and Gordon (2011)). I use benchmarks from the PerformanceTEST software developed by PassMark Software (2026a). One complication is that multiple VMs may share a single CPU, or even a single core. Some VM instances have a single vCPU, which corresponds to a single thread (Amazon Web Services (nd)). Thus, the total multithread performance on PassMark Software (2026a) may not adequately represent VM quality. To compute the relative performance of a VM, therefore, I take the multithread score and divide it by the number of threads on a CPU to get a per-thread performance measure. For example, Microsoft Azure's A10 instance uses Intel Xeon E5-2660 v0, which has a multithread score of 8076 and a total of 16 threads, for a per-thread score of 504.75 (PassMark Software (2026b)). Wherever possible, I use official documentation; otherwise, I use third party cloud platform comparison websites, which run benchmarking software that is able to identify the underlying processor of a given instance type. When multiple CPU products are used for the same VM instance type, I take the average score. Finally, I follow the same approach as in Appendix A.6.1 to compute the average benchmark score at the provider-datacenter-year level. Table A.5 lists my sources for the CPUs used by each VM instance.

Source Name	URL
Official Documentation	
Azure A-Series & Av2-Series Virtual Machines	docs.microsoft.com/.../av2-series
Azure Bsv2 / Bs-Series Virtual Machines	docs.microsoft.com/.../bsv2-series
Azure Dv2 and DSv2-Series Virtual Machines	docs.microsoft.com/.../dv2-dsv2-series
Azure Dv3 and Dsv3-Series Virtual Machines	docs.microsoft.com/.../dv3-dsv3-series
Azure Dv4 and Dsv4-Series Virtual Machines	docs.microsoft.com/.../dv4-dsv4-series
Azure Dv5 and Dsv5-Series Virtual Machines	docs.microsoft.com/.../dv5-dsv5-series
Azure Dav4 and Dasv4-Series Virtual Machines	docs.microsoft.com/.../dav4-dasv4-series
Azure Dasv5 and Dadsv5-Series Virtual Machines	docs.microsoft.com/.../dasv5-dadsv5-series
Azure DCsv2 / DC-Series (Confidential Computing)	docs.microsoft.com/.../dcv2-series
Azure Ev3 and Esv3-Series Virtual Machines	docs.microsoft.com/.../ev3-esv3-series
Azure Ev4 and Esv4-Series Virtual Machines	docs.microsoft.com/.../ev4-esv4-series
Azure Ev5 and Esv5-Series Virtual Machines	docs.microsoft.com/.../ev5-esv5-series
Azure Eav4 and Easv4-Series Virtual Machines	docs.microsoft.com/.../eav4-easv4-series
Azure Easv5 and Eadsv5-Series Virtual Machines	docs.microsoft.com/.../easv5-eadsv5-series
Azure F-Series, Fs-Series, and Fsv2-Series VMs	docs.microsoft.com/.../fsv2-series

(continued on next page)

(continued from previous page)

Source Name	URL
Azure G-Series and Gs-Series Virtual Machines	docs.microsoft.com/.../sizes-previous-gen
Azure H-Series Virtual Machines	docs.microsoft.com/.../h-series
Azure Ls-Series and Lsv2-Series Virtual Machines	docs.microsoft.com/.../lsv2-series
Azure NC-Series Virtual Machines	docs.microsoft.com/.../nc-series
Azure NCas T4 v3-Series Virtual Machines	docs.microsoft.com/.../nct4-v3-series
Azure NV-Series Virtual Machines	docs.microsoft.com/.../nv-series
Azure NVv4-Series Virtual Machines	docs.microsoft.com/.../nvv4-series
Azure: Sizes for Virtual Machines (Overview)	docs.microsoft.com/.../sizes
Azure: Previous Generation VM Sizes	docs.microsoft.com/.../sizes-previous-gen
GCP: N1 General-Purpose Machine Family	cloud.google.com/.../general-purpose-machines#n1
GCP: N2 General-Purpose Machine Family	cloud.google.com/.../general-purpose-machines#n2
GCP: N2D General-Purpose Machine Family	cloud.google.com/.../general-purpose-machines#n2d
GCP: E2 General-Purpose Machine Family	cloud.google.com/.../general-purpose-machines#e2
GCP: C2 Compute-Optimized Machine Family	cloud.google.com/.../compute-optimized-machines#c2
GCP: T2D General-Purpose Machine Family	cloud.google.com/.../general-purpose-machines#t2d
GCP: Machine Families Resource and Comparison Guide	cloud.google.com/.../machine-resource
Azure Ebdsv5 and Ebsv5-Series Virtual Machines	learn.microsoft.com/.../ebdsv5-ebsv5-series
Third-Party Sources	

(continued on next page)

(continued from previous page)

Source Name	URL
PassMark: ARM Cortex-A72 16 Core 2000 MHz	cpubenchmark.net/.../id=6436
PassMark: ARM Neoverse-N1 4 Core	cpubenchmark.net/.../id=4397
Geekbench 2: Intel Xeon E5506 @ 2.13 GHz (EC2)	browser.geekbench.com/geekbench2/compare/1585176
Geekbench 3: Intel Xeon E5507 @ 2.27 GHz (EC2)	browser.geekbench.com/v3/cpu/546695
Geekbench 5: Amazon EC2 t2.2xlarge	browser.geekbench.com/v5/cpu/16843977
Geekbench 6: AWS EC2 t2.large (baseline)	browser.geekbench.com/v6/cpu/compare/2502090
Geekbench 6: Amazon EC2 t3.2xlarge vs t2.large	browser.geekbench.com/v6/cpu/compare/9526213
Geekbench 6: Amazon EC2 t3a.large vs t2.large	browser.geekbench.com/v6/cpu/compare/3064894
VPSBenchmarks: Google Compute Engine CPUs	vpsbenchmarks.com/.../google_compute_engine/cpus
gcloud-compute.com: f1-micro Specifications	gcloud-compute.com/f1-micro.html
emma: AWS EC2 g4dn.xlarge Spec & Pricing Guide	emma.ms/blog/aws-ec2-g4dn-xlarge
CloudOptimo: AWS EC2 m1.large Pricing & Specs	costcalc.cloudoptimo.com/.../m1.large
RunsOn: AWS EC2 m6g Instance Family	go.runs-on.com/instances/ec2/m6g
EC2Instances.info: Easy Amazon EC2 Instance Comparison	ec2instances.info

A.6.4 Datacenter locations

The map in Figure A.3 provides locations of datacenters for each of the cloud providers in my data at the beginning and end of each sample period (Datacenters.com (2025)).

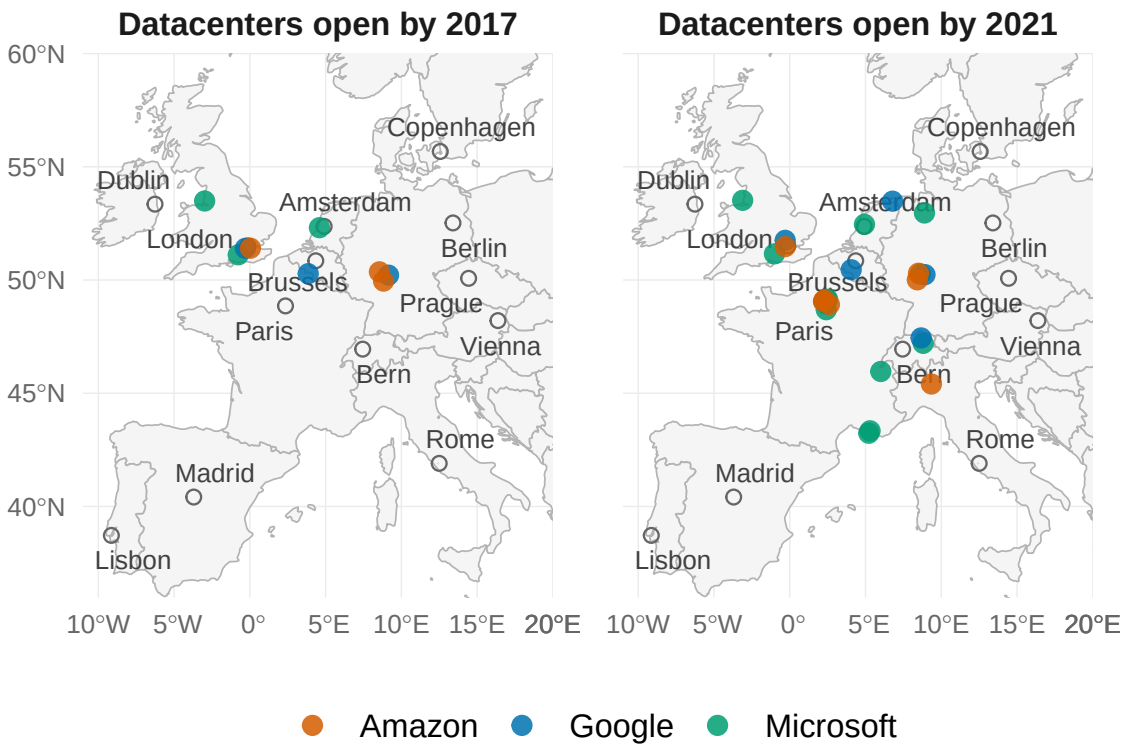


Figure A.3: Datacenter locations in 2017 and 2021

A.6.5 Feature and region merge

Since it is far more complete than the Orbis location data, I use the headquarter latitude and longitude data from Aberdeen to match firms to datacenters. The matching algorithm is as follows. For each firm, I compute the Haversine distance between the firm's headquarters and the nearest datacenter for each provider used, where headquarters are defined as follows. If the firm is single-establishment, I use that establishment. If the firm has an establishment designated as "Ultimate HQ," which is unique in the Aberdeen data, I use that. Otherwise, if there is at least one establishment designated as "Headquarters," I randomly choose one. Finally, if the firm has no in-country headquarters, I randomly assign an establishment labeled as "Branch." I drop firms that do not have location coordinate data. I assume that firms choose a datacenter upon adopting a cloud computing service and never switch within-supplier datacenters, consistent with prior research showing that such changes are costly and therefore rare (Demirer et al. (2024)).

As this process may introduce measurement error in prices and qualities, I compute summary statistics for 1) my full matched sample, and 2) my matched sample excluding matches on "Branch" establishments, and 3) my matched sample excluding matches on "Branch" and "Headquarter" establishments.[85] The results are presented in Table A.6.

Table A.6: Robustness to excluding multinational firms

	$\bar{p}_{Amazon}$	$\bar{p}_{Google}$	$\bar{p}_{Microsoft}$
Full sample	0.051296	0.047661	0.059263
Excluding "Branch" matches	0.050237	0.047697	0.059290
Excluding "HQ" and "Branch" matches	0.051639	0.047814	0.059933

Means taken across all downstream firms with location data, including set of potential entrants. Prices in real 2015 euros.

A.7 Microsoft server products

A full list of Microsoft server-related products is displayed in Table A.7.

[85] Even if multinational firms have a local headquarters in France, they still may store data at the datacenter nearest their global headquarters. I thank Daniela Muhaj for making this observation.

Table A.7: List of Microsoft server products

Category	Product name
Windows Server	Windows NT Server; Windows Server; Microsoft Virtual Server; Hyper-V Server
SQL Server	Microsoft SQL Server ; SQL Server Reporting Services; SQL Server Analysis Services; SQL Server Management Studio; SQL Server Business Intelligence; Transact-SQL
Exchange Server	Microsoft Exchange Server; Exchange Archiving; Exchange Online Protection
SharePoint	Microsoft SharePoint; Microsoft Enterprise Content Management
Unified Communications	Microsoft Lync; Lync Server; Lync Enterprise Voice; Skype for Business
System Center	Microsoft System Center; System Center Operations Manager; Configuration Manager; Data Protection Manager; Endpoint Protection; Virtual Machine Manager; Server Management Tools
Identity and Directory Services	Active Directory Federation Services; Kerberos Authentication; Group Policy Management; Microsoft DirectAccess
Virtualization and Application Virtualization	Microsoft Hyper-V; Microsoft Application Virtualization (App-V)
Enterprise Integration and Middleware	Microsoft BizTalk Server; Microsoft Commerce Server; Enterprise Project Management; Customer Care Framework; Enterprise Application Integration
Hosting, Storage, and Mobility Infrastructure	Microsoft Hosting Services; Web Hosting; Microsoft StorSimple; Storage Management; Microsoft Enterprise Mobility Suite
Security and Threat Management	Microsoft Forefront; Forefront Unified Access Gateway

List of all product names that have appeared it at least one year of Aberdeen's data from 2015-2021. Names have been standardized across survey waves.

A.8 Calibrating marginal costs

This appendix contains details for calibrating the marginal costs of supplying a vCPU for one hour and for the network costs of moving one GB of data off platform.

A.8.1 vCPU marginal cost

To model the marginal cost of running a vCPU for one hour, I rely on prior engineering estimates of the variable costs of running cloud computing servers (Mach and Schikuta (2011)). I can write the cost model as follows:

$$
\begin{aligned}
C_{jt}^{vm} &= C_{jt}^{depr} + C_{jt}^{power} \\
\text{where } C_{jt}^{depr} &= \frac{\delta^{VM}}{8760} \cdot \frac{\bar{C}_{jt}^{server}}{\#vCPU} \\
\text{and } C_{jt}^{power} &= (0.5 + 0.5\bar{U}_{jt})\bar{\varkappa}_{jt}^{max} \cdot P_{jt}^{elec}
\end{aligned}
$$

where δ^{VM} is the annual depreciation rate for computing hardware, calibrated to be 0.36 (Bloom et al. (2012)), which I divided by 8760 to convert to hourly depreciation. Then, $\bar{C}_{jt}^{server}, \bar{C}_{t}^{SSD}$ are the mean server and SSD prices, weighted by usage as described above. I do not observe oligopsony power in the input market, so I instead use the resale price of CPUs from Passmark.com as a proxy for the cost. The expression for C_{jt}^{power} comes from Brand et al. (2024), where $\bar{U}_{jt}$ is the average annual usage rate; $\bar{\varkappa}_{jt}^{max}$ is the mean maximum energy consumption, which I proxy using manufacturers' thermal density power (TDP) data from Passmark.com; and P_{jt}^{elec} is the average annual electricity price in kWh, weighted by consumer-level datacenter choices. Electricity prices come from Eurostat (France, Germany, Netherlands, Belgium); DESNZ (United Kingdom); and ElCom (Switzerland). All prices are converted to euros; for each country, I use the largest band grid. In the model, calibrated marginal costs are weighted by downstream firms' choice of datacenter. Results are displayed in Table A.8.

Table A.8: Marginal cost summary statistics

	Amazon	Google	Microsoft
Marginal cost	0.00132 (0.000837)	0.00205 (0.000334)	0.00135 (0.000354)

Mean costs in real euros per vCPU hour across datacenter-years are displayed, with standard deviations in parentheses.

A.8.2 Egress marginal cost

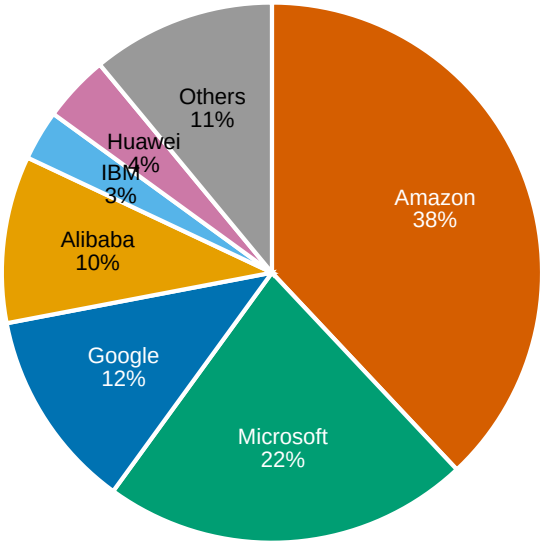
The main component of egress fees are the network costs of transferring data. I use the lowest competitive rate in 2022 from Boudreau (2022) of \$0.05/Mbps-month, or approximately 0.00015 euros per GB.

$$C^{egress} = \frac{0.05\text{USD}}{\text{Mbps-month}} \cdot \frac{1}{324} \approx 0.00015 \frac{EUR}{GB}$$

This is near the findings of France's telecommunications regulatory body, ARCEP, which finds the marginal cost of data egress to be €0 (Arcep (2025)).

A.9 Evidence of worldwide IaaS market concentration

Figure A.4 provides external evidence that cloud computing users tend to use a primary provider by showing the deciles of the data distribution for the customers of Wiz, a large cloud security firm.



Source: Statista (2026)

Figure A.4: Worldwide market share of IaaS firms (2023)

A.10 Difference-in-differences regression on cloud spot prices

To provide evidence about whether or not egress costs matter for firm switching in the short run, I leverage the random timing of the major firms' decisions to drop egress costs on the spot market price of virtual machines. In particular, Google announced it would drop fees for consumers exiting their platform on January 11, 2024, while Amazon and Microsoft announced their fee eliminations on March 5 and 13, 2024, respectively To estimate the short run effects of an egress fee elimination on demand for virtual machines, I run a difference-in-differences specification for the 54 days surrounding Google's announcement. In the short run, if egress fees are a barrier to switching, one should expect to see a fall in the spot market price relative to the list price as consumers exit. I compute the discount offered by the spot market compared to the list price for virtual machine instance p on day d as:

$$\text{Savings}_{pd} = \frac{P_{pd}^{spot}}{P_{pd}^{list}} \times 100\% \tag{32}$$

I use Google's announcement as a treatment and run the following difference-in-differences specification:

Variable	Model 1 (54 days)	Model 3 (30 days)
Treatment	-0.0255 (0.0243)	0.0262 (0.0235)
N. obs	4,639,571	2,574,119
Adj. R^2	0.313944	0.624046
Product FE	✓	✓
Supplier FE	✓	✓
Region FE	✓	✓
Day FE	✓	✓

Table A.9: Effect of egress fee elimination on VM spot prices

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run with product, provider, time, and region fixed effects and use heteroskedasticity-robust standard errors.

$$\text{Savings}_{pd} = \sum_{d \neq -1} \beta \cdot \mathbb{1}\{j = \text{Google}\} \cdot \mathbb{1}\{t \geqslant \text{Jan 11, 2024}\} + \alpha_p + \beta_j + \tau_d + \gamma_r + \epsilon_{pd} \quad (33)$$

where $\mathbb{1}\{j = \text{Google}\}$ is an indicator; and $\alpha_p, \beta_j, \tau_d$, and γ_r are product, supplier, region r, and time fixed effects, respectively. I use the companion dataset from Lee et al. (2022) and average spot price discounts at the daily level.[86] I run a specification for the full data ($\pm$ 54 days around treatment) as well as a shorter window ($\pm$ 30 days around treatment). Results are presented in Table A.9 below. Both regressions produce null results.

B Modeling details appendix

This appendix derives additional results for the structural model. Appendix B.1 derives the expression for a downstream firm's capital investment choice, following Gowrisankaran and Rysman (2012). Appendix B.2 derives the closed form expressions for firms' IT bundle choices. Using the derivations in the prior two appendices, Appendix B.3 derives the suppliers' pricing first order conditions in closed form, which facilitates model convergence and accelerates estimation. Finally, Appendix B.4 microfounds the relationship between a firm's overall capital stock, K_{ft}, and its data stock, D_{ft}, using prior literature (Farboodi and Veldkamp (2021)).

[86] The data for Amazon contains data at the region and availability zone (AZ) level. Since AZ names are anonymized for each customer and I only observe data for a single customer, I also average across AZ.

B.1 Capital investment details

Following Yun (2026); Gowrisankaran and Schmidt-Dengler (2024) and suppressing dependence on sector i, let the choice-specific value for selecting some capital stock K^j be

$$v_{ft}^j = -C^K(K^j, K_{ft}) + \beta \mathbb{E}_t[V_{t+1}(s_{ft+1}, \psi_{t+1})|s_{ft}, \psi_t] \tag{34}$$

for the j-th grid of the set of capital values, $\mathcal{K}$. Given a bundle choice b_{ft}, a firm chooses K^j if and only if

$$V_{ft}^j - I_{ft}^j \epsilon_{ft}^K \geqslant V_{ft}^k - I_{ft}^k \epsilon_{ft}^K \tag{35}$$

where

$$I_{ft}^j = K^j - (1-\delta)K_{ft} \tag{36}$$

Monotonicity of $I(\cdot, K_{ft})$ implies that the firm prefers K^j over K^k if and only if

$$\epsilon_{ft}^K \geqslant \frac{v_{ft}^k - v_{ft}^j}{I_{ft}^k - I_{ft}^j} \equiv \epsilon_{ft}^K(j,k) \tag{37}$$

for any $j > k$. The firm chooses K^j if and only if $\underline{\epsilon}_{ft}(j) \leqslant \epsilon_{ft}^K < \bar{\epsilon}_{ft}(j)$, where

$$\underline{\epsilon}_{ft}(j) = \begin{cases} \epsilon_{ft}^K(j, j+1), & j < J \\ -\infty, & j = J \end{cases} \tag{38}$$

$$\bar{\epsilon}_{ft}(j) = \begin{cases} \infty, & j = 1 \\ \epsilon_{ft}^K(j-1, j), & j > 1 \end{cases} \tag{39}$$

Then the conditional probability of choosing capacity K^j in state s_{ft} is given by

$$\mathbb{P}[K_{ft} = K^j | s_{ft}, \psi_t] = \Phi\left(\frac{\bar{\epsilon}_{ft}(j)}{\sigma^K}\right) - \Phi\left(\frac{\underline{\epsilon}_{ft}(j)}{\sigma^K}\right) \tag{40}$$

Finally, I can follow Gowrisankaran and Schmidt-Dengler (2024) and use the properties of the truncated normal distribution to derive the conditional mean of the shock to be

$$\mathbb{E}[\epsilon_{ft}|K_{ft} = K^j, s_{ft}, \psi_t] = \frac{\phi\left(\frac{\underline{\epsilon}_{ft}(j)}{\sigma^K}\right) - \phi\left(\frac{\bar{\epsilon}_{ft}(j)}{\sigma^K}\right)}{\Phi\left(\frac{\bar{\epsilon}_{ft}(j)}{\sigma^K}\right) - \Phi\left(\frac{\underline{\epsilon}_{ft}(j)}{\sigma^K}\right)} \tag{41}$$

B.2 Bundle choice details

The functional form of the entry and maintenance cost shocks means I can rewrite firms' value functions as follows. Throughout, I suppress dependence on industry i.

First, denote

$$\sigma^E(K_{ft}) \equiv \sigma^E \cdot D_{ft}(K_{ft}) \tag{42}$$

$$\sigma^b(K_{ft}) \equiv \sigma^b \cdot D_{ft}(K_{ft}) \tag{43}$$

as the effective cost shock incurred by a firm of size K_{ft}. Denote the value function at state s_{ft} conditional on $b;, K'$ as follows:

$$v_t(s_{ft}, b', K') = \pi_t(s_{ft}) - C^K(I(K'), K_{ft}, \epsilon^K_{ft}) - D_{ft}\kappa_{ft}(s_{ft}) + \beta \mathbb{E}_t[V_{t+1}(s_{ft+1})|s_{ft}]$$

where

$$C^K(I(K'), K_{ft}, \epsilon^K_{ft}) = p^K_{it} I(K') + C^K_1 \left(\frac{I(K')^2}{K_{ft}} \right) K_{ft} + I(K')\epsilon^K_{ft}$$

$$\kappa_{ft}(s_{ft}) = \left\{ \begin{array}{l} \kappa^E_t \cdot \mathbb{1}\{b_{ft+1} \neq b_{ft} | b_{ft} = \varnothing\} + \kappa^S \cdot \mathbb{1}\{b_{ft+1} \neq b_{ft} | b_{ft} \neq \varnothing\} + \\ \kappa^M(\boldsymbol{\alpha}' \mathbf{X}_{bft}) \cdot \mathbb{1}\{b_{ft} \neq \varnothing\} \end{array} \right\}$$

Denoting the value function minus the investment cost shock as

$$\tilde{v}_t(s_{ft}, b', K') = v_t(s_{ft}, b', K') - I(K')\epsilon^K_{ft}$$

Using the results from B.1, I can write the investment-smoothed version of the value function as

$$\hat{v}_t(s_{ft}, b') = \sum_j \left[\Phi\left(\frac{\bar{\epsilon}_{ft}(j)}{\sigma^K} \right) - \Phi\left(\frac{\underline{\epsilon}_{ft}(j)}{\sigma^K} \right) \right] \cdot \left(\tilde{v}_t(s_{ft}, b', K^{j\prime}) - I^j \mathbb{E}[\epsilon_{ft} | K_{ft} = K^j, s_{ft}] \right)$$

for all $K^j \in \mathcal{K}$. Let $\mathcal{B}^+ = \mathcal{B} \backslash \{\varnothing\}$. Then, using the assumptions that $\epsilon^E_{ft} \sim Gumbel(0, \sigma^E)$, $\epsilon^b_{ft} \sim Gumbel(0, \sigma^b)$ I can write the value functions for non-adopters, existing consumers, and new

entrants as follows:

$$
\begin{aligned}
V_t^0(s_{ft}) &= \hat{v}_t(s_{ft}, b') \\
V_t^{Inc}(s_{ft}) &= \sigma^b(K_{ft}) \log \sum_{b' \in \mathcal{B}^+} \exp\left(\frac{\hat{v}_t(s_{ft}, b')}{\sigma^b(K_{ft})}\right) + \sigma^b(K_{ft}) \cdot \tilde{\gamma} \\
V_t^{Ent}(s_{ft}) &= \sigma^E(K_{ft}) \log\left[\exp(V_t^0(s_{ft})/\sigma^E(K_{ft})) + \exp(V_t^{Inc}(s_{ft})/\sigma^E(K_{ft}))\right] + \sigma^E(K_{ft}) \cdot \tilde{\gamma}
\end{aligned}
$$

where $\tilde{\gamma}$ is the Euler-Mascheroni constant. Finally, I can write the choice probabilities for entry and for the conditional bundle choice as follows:

$$
\begin{aligned}
\varsigma_t^{Entry}(s_{ft}) &\equiv \mathbb{P}_t[b_{ft+1} \neq \varnothing | b_{ft} = \varnothing, s_{ft}] = \frac{\exp(V_t^{Inc}(s_{ft})/\sigma^E(K_{ft}))}{\exp(V_t^0(s_{ft})/\varsigma^E(K_{ft})) + \exp(V_t^{Inc}(s_{ft})/\sigma^E(K_{ft}))} \\
\varsigma_t^{Bundle}(s_{ft}) &\equiv \mathbb{P}_t[b_{ft+1} = b' | b_{ft} \neq \varnothing, s_{ft}] = \frac{\exp(\tilde{v}_t(s_{ft}, b', K')/\sigma^b(K_{ft}))}{\sum_{\hat{b}' \in \mathcal{B}} \exp(\tilde{v}_t(s_{ft}, \hat{b}', K')/\sigma^b(K_{ft}))}
\end{aligned}
$$

B.3 Supplier pricing details

In this section, I derive suppliers' value functions and first order conditions. I use the derivations in the above appendices to write closed form versions of each. Before downstream firms make their investment choices in each period, upstream firms choose prices P_{jt} to maximize their value function:

$$
\begin{aligned}
W_{jt} &= \max_{P_{jt}} \mathbb{E}_t[\pi_{jt}^S + \beta W_{jt+1}(s_{t+1}) | s_t] \\
\text{where } \pi_{jt}^S &= (P_{jt} - C_{jt}^{vm}) Q_{jt} + (E_{jt} - C_j^{egress})[Q_{jt}^{exit} + \lambda Q_{jt}^{multi}]
\end{aligned}
$$

for $\lambda \in [0, 1]$. I assume that E_{jt} is set exogenously and $w_{jb} = \frac{\mathbb{1}\{j \in b\}}{\#j \in b}$. Let the strategy of cloud provider j in year t be given by $P_{jt} \in \mathbb{R}$. Let $\tilde{s}_{ft} = s_{ft} \backslash \mathbf{P}_t, b_{ft}$. Finally, N_{it}^b is the mass of firms in industry i using bundle b in period t and $dF_{it}(\tilde{s}_{ft} | b_{ft})$ is the industry-year distribution of $\tilde{s}_{ft}$ given b_{ft}. Using this, I take the first order conditions with respect to P_{jt} as

$$
Q_{jt} + (P_{jt} - C_{jt}^{vm}) \cdot \frac{\partial Q_{jt}}{\partial P_{jt}} + (E_{jt} - C_j^{egress}) \cdot \left[\lambda \frac{\partial Q_{jt}^{multi}}{\partial P_{jt}} + \frac{\partial Q_{jt}^{exit}}{\partial P_{jt}}\right] + \beta \frac{\partial \mathbb{E}_t W_{jt+1}(s_{t+1})}{\partial P_{jt}} = 0
$$

Next, using the conditional choice probabilities derived above in Appendix B.2, I write provider j's quantities at period t as

$$Q_{jt} = \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}} N_{it}^{b} \int \varsigma_t^{Entry}(\tilde{s}_{ft}, b') \sum_{b'\ni j} w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') \alpha_1 D_{ft} \; dF_{it}(\tilde{s}_{ft}|b_{ft})$$

$$Q_{jt}^{multi} = \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}} N_{it}^{b} \int \varsigma_t^{Entry}(\tilde{s}_{ft}, b') \sum_{b'\ni j} \mathbb{1}\{\#j_{b'} > 1\} \cdot w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') D_{ft} \; dF_{it}(\tilde{s}_{ft}|b_{ft})$$

$$Q_{jt}^{exit} = \begin{cases} \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}} N_{it}^{b} \int (1 - \varsigma_t^{Entry}(\tilde{s}_{ft}, b')) \sum_{b'\ni j} w_{jb'} \alpha_1 D_{ft} \; dF_{it}(\tilde{s}_{ft}|b_{ft}) & \text{if } b' = \varnothing \\ \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}} N_{it}^{b} \int \varsigma_t^{Entry}(\tilde{s}_{ft}, b') \sum_{b'\ni j} [w_{jb'} - w_{jb}]_+ \cdot \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') D_{ft} \; dF_{it}(\tilde{s}_{ft}|b_{ft}) & \text{if } b' \neq \varnothing \end{cases}$$

The derivatives of each expression are as follows. For Q_{jt}:

$$\frac{\partial Q_{jt}}{\partial P_{jt}} = -\alpha_! \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}} N_{it}^{b} \int \left[\frac{\varsigma_t^{Entry}(\tilde{s}_{ft}, b')(1 - \varsigma_t^{Entry}(\tilde{s}_{ft}, b'))}{\sigma^E(K_{ft})} \left(\sum_{b'\ni j} w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') \alpha_1 D_{ft} \right)^2 \right.$$
$$\left. + \frac{\varsigma_t^{Entry}(\tilde{s}_{ft}, b')}{\sigma^b(K_{ft})} \sum_{b'\ni j} w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') \alpha_1 D_{ft} \left(w_{jb'} \alpha_1 D_{ft} - \sum_{b'\ni j} w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') \alpha_1 D_{ft} \right) \right] dF_{it}(\tilde{s}_{ft}|b_{ft})$$

For Q_{jt}^{multi}:

$$\frac{\partial Q_{jt}^{multi}}{\partial P_{jt}} = -\alpha_1 \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}} N_{it}^{b} \int \left[\frac{\varsigma_t^{Entry}(\tilde{s}_{ft}, b')(1 - \varsigma_t^{Entry}(\tilde{s}_{ft}, b'))}{\sigma^E(K_{ft})} \times \right.$$
$$\left(\sum_{b'\ni j} w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') D_{ft} \right) \left(\sum_{b'\ni j} \mathbb{1}\{\#j_{b'} > 1\} \cdot w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') D_{ft} \right)$$
$$+ \frac{\varsigma_t^{Entry}(\tilde{s}_{ft}, b')}{\sigma^b(K_{ft})} \sum_{b'\ni j} \mathbb{1}\{\#j_{b'} > 1\} \cdot w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') D_{ft} \times$$
$$\left. \left(w_{jb'} D_{ft} - \sum_{b'\ni j} w_{jb'} \varsigma_t^{Bundle}(\tilde{s}_{ft}, b') D_{ft} \right) \right] dF_{it}(\tilde{s}_{ft}|b_{ft})$$

Finally, for Q_{jt}^{exit}, I have separate derivatives for whether a firm ceases to use cloud computing entirely or moves data out of a (sub)set of suppliers to another cloud bundle. First, for complete exit, the expression is as follows:

$$\frac{\partial Q_{jt}^{exit}}{\partial P_{jt}} = \alpha_1 \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}} N_{it}^b \int \frac{\varsigma_t^{Entry}(\tilde{s}_{ft},b')(1-\varsigma_t^{Entry}(\tilde{s}_{ft},b'))}{\sigma^E(K_{ft})} \left(\sum_{b'\ni j} w_{jb'}\varsigma_t^{Bundle}(\tilde{s}_{ft},b')D_{ft}\right) w_{jb'}\ dF_{it}(\tilde{s}_{ft}|b_{ft})$$

Second, for exit to another cloud bundle, it is:

$$\begin{aligned}\frac{\partial Q_{jt}^{exit}}{\partial P_{jt}} = -\alpha_1 \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}} N_{it}^b \int \Bigg[& \frac{\varsigma_t^{Entry}(\tilde{s}_{ft},b')(1-\varsigma_t^{Entry}(\tilde{s}_{ft},b'))}{\sigma^E(K_{ft})} \times \\ & \left(\sum_{b'\ni j} w_{jb'}\varsigma_t^{Bundle}(\tilde{s}_{ft},b')D_{ft}\right) \sum_{b'\notin\varnothing} \varsigma_t^{Bundle}(\tilde{s}_{ft},b')[w_{jb}-w_{jb'}]_+ \\ & + \frac{\varsigma_t^{Entry}(\tilde{s}_{ft},b')}{\sigma^b(K_{ft})} \sum_{b'\ni j} \mathbb{1}\{\#j_{b'}>1\}\cdot w_{jb'}\varsigma_t^{Bundle}(\tilde{s}_{ft},b')D_{ft} \times \\ & \left(w_{jb'}D_{ft} - \sum_{b'\ni j} w_{jb'}\varsigma_t^{Bundle}(\tilde{s}_{ft},b')D_{ft}\right)\Bigg]\ dF_{it}(\tilde{s}_{ft}|b_{ft})\end{aligned}$$

Note that because suppliers internalize some of the complementarities of multihoming through charging egress fees on per-period data flows, this system of equations does not always satisfy supermodularity for all parameter values and thus does not guarantee a unique fixed point. While I find empirically that the system converges to a fixed point near the true parameter values, I describe in Appendix C.3 how I deal with convergence failures far from the global optimum.[87]

B.4 A model of capital and data

In this appendix, I show how I recover the parameters governing firms' data stocks as a function of their capital stock. Suppose that, in each industry i, production is Leontief in data and non-IT capital:

$$K_{ft} = \min\{\varrho_0 D_{ft}^{\varrho_1}, \tilde{K}_{ft}e^{\epsilon_{ft}^D}\}$$

where

$$\tilde{K}_{ft} = K_{ft} - D_{ft}$$

and $e^{\epsilon_{ft}^D}$ is a multiplicative error term. This can be rationalized as each machine or other

[87] To provide intuition, firms pay $P_{jt}\cdot\alpha_1 D_{ft}$ to use vCPUs, but only $E_{ft}D_{ft}$ to egress data. For large α_1, the former dominates pricing incentives.

unit of capital requiring a certain fraction of data to operate.[88] Therefore, I can write

$$D_{ft} = \left(\frac{1}{\varrho_0} \tilde{K}_{ft} e^{\epsilon^D_{ft}} \right)^{1/\varrho_1}$$

Taking logs, I can write

$$\begin{aligned} \log(D_{ft}) &= \upsilon_{0i} + \upsilon_{1i} \log(K_{ft}) + \Xi_\iota + \tau_t + \epsilon^D_{ft} \\ \text{where } \upsilon_{0i} &= \left(\frac{1}{\varrho_0} \right)^{1/\varrho_1} \\ \upsilon_{1i} &= 1/\varrho_1 \end{aligned}$$

Given my limited pre-cloud computing era data (2015-2016), I estimate this using OLS, where Ξ_ι are 4-digit NACE industry fixed effects and τ_t is a year fixed effect.

C Estimation appendix

In this appendix I provide further details on how I estimate parameters. Estimation of the production function is standard and is omitted. First, I describe the estimation of the transition process for exogenous states (C.1). Second, I describe estimation of my remaining parameters, detailing the MME algorithm and MSM procedure (C.2). Finally, I expand on how the MSM parameters are computed (C.3).

C.1 Exogenous state transitions

This appendix provides details on how I model transitions of the set of exogenous state variables.

Microsoft server usage. To model the evolution of Microsoft server product adoption, I run the following AR(1) model for each sector i:

$$server_{ft} = \rho^s_0 + \rho^s_i server_{ft-1} + \epsilon^s_{ft}$$

The results are displayed in Table C.1.

Virtual machine quality. I assume that, beyond the data years, log VM quality follows

[88] Lashkari et al. (2024) find a low estimate of the substitution elasticity between IT inputs and non-IT inputs, though greater than 0.

Table C.1: AR(1) persistence of Microsoft server adoption (2017-2021)

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Intercept	0.0735 (0.0018)***	0.0453 (0.0008)***	0.0389 (0.0003)***
ρ_i^s	0.8917 (0.0033)***	0.9177 (0.0019)***	0.9238 (0.0007)***
Observations	25,328	82,962	582,236
Residual Std. Error	0.2397	0.2049	0.1927
R^2	0.7470	0.7311	0.7603
Adj. R^2	0.7470	0.7311	0.7603

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Note: OLS regressions of Microsoft server adoption on its lag.

an exogenous exponential growth rate, as a consequence of Moore's law.[89][90] The regression for log quality has $N = 61$ observations and is given by

$$d\log(q_{jt}) = \underbrace{\rho^q}_{0.0322(0.009)} + \epsilon_{jt}^q$$

Other aggregate state variables. I have three other time-varying state variables: mean industry wage, mean industry materials price, and the capital rental rate. As they are not a focus of my analysis, I use industry-level average values for $t \in [2017, 2021]$ and then assume that the industry-level wage, materials price, and rental rate remain fixed at their 2021 values for all $t > 2021$.

C.2 Dynamic parameter estimation details

In this appendix, I provide details on estimating the dynamic choice parameters. I use the Markov moment equilibrium (MME) (Ifrach and Weintraub (2017)) approximation to an MPE for supplier prices.[91] I estimate the parameters via the method of simulated moments (MSM) (McFadden (1989); Pakes and Pollard (1989)).

Algorithm. The algorithm for computing an MME is similar to standard value function iteration for consumers and Bertrand pricing for suppliers, but with an added outer loop. First, I guess a candidate transition path of prices $\hat{M}^\psi$. In the inner loop, I solve the steady state value functions for consumers at some period T in the future, using prices at p_T as an

[89] I also run a version with cloud provider fixed effects, but these are not statistically significant.

[90] Quality does not perfectly follow the rate of improvements in semiconductors because firms continue to offer VMs based on old vintages.

[91] Accounting for nonstationarity closely follows Benkard et al. (2024), which extends the oblivious equilibrium concept to study short-run dynamics from common shocks.

approximation. I then compute values along the transition path using backwards induction. Finally, I simulate prices forward and solve for consumers' best responses and suppliers' Bertrand pricing strategies. I then fit an AR(1) model over the path of each supplier j's price path, $\hat{M}^{\psi}$. The algorithm repeats until convergence of $\hat{M}^{\psi}$.

Formally, denote decision rules for downstream firms and suppliers as follows:

$$
\begin{aligned}
a_{ft} &= \psi_t^d(s_{ft}, \sigma_{ft}^K, \sigma_{ft}^E, \sigma_{ft}^b) \\
P_{jt} &= \psi_t^s(s_t)
\end{aligned}
$$

such that the year-t strategy profile can be written as

$$
\psi^t = \{\psi_t^d(s_{ft}, \sigma_{ft}^K, \sigma_{ft}^E, \sigma_{ft}^b)\}_{f \in \mathcal{F}} \cup \{\psi_t^s(s_t)\}_{j \in \{1,2,3\}}
$$

Then the perceived transition kernel for the is given by

$$
\hat{M}_t^{\psi} = \Phi M_t^{\psi} \tag{44}
$$

where $\hat{M}_t^{\psi}$ is the transition kernel of the moment-based state when agents use strategy ψ_t, M_t^{ψ} is the true transition kernel, and Φ is an operator such that Equation 44 holds. I approximate M_t^{ψ} as

$$
\begin{bmatrix} P_{Amazon,t+1} \\ P_{Google,t+1} \\ P_{Microsoft,t+1} \end{bmatrix} = \begin{bmatrix} \rho_0^{Amazon} \\ \rho_0^{Google} \\ \rho_0^{Microsoft} \end{bmatrix} + \begin{bmatrix} \rho_0^{Amazon} P_{Amazon,t} & 0 & 0 \\ 0 & \rho_0^{Google} P_{Google,t} & 0 \\ 0 & 0 & \rho_0^{Microsoft} P_{Microsoft,t} \end{bmatrix}
$$

The MME algorithm is detailed below. For each candidate parameter vector $\hat{\boldsymbol{\theta}}$, run:

1. Initialize $V_{ft}(s_{ft})$ for all $s_{ft} \in \mathcal{S}$.
2. Run inner loop:
   (a) Solve for terminal period value functions $V_T(s_{fT})$ of downstream firms at the terminal T for all states s_{fT}.
   (b) Compute value functions along the transition path $t = 0, .., T-1$ by backwards induction.
   (c) Initialize the distribution of states in $t = 0$, s_0, using data.

(d) Simulate the industry forward and compute supplier and downstream firm decision rules at each state to obtain actions ψ^t, forecast prices $\{\hat{P}_{jt}\}_{j\in\{1,2,3\},t\in\mathcal{T}}$, and therefore $\hat{M}_\psi$.

3. Calculate the updated value functions $\tilde{V}_{ft}(s_{ft})$ and new equilibrium decision rules ψ^t.

4. Re-run the inner loop and obtain an updated $\hat{M}_\psi$.

5. Compute the norm $||\{\frac{\hat{M}_\psi - M_\psi}{M_\psi}\}||$, where $||\cdot||$ is the sup norm.

6. If the norm is greater than $\epsilon^{tol} = 0.001$, repeat steps 2-5. Else terminate.

Optimization. Let $\boldsymbol{\theta} = [\kappa_0^E, \kappa_1^E, \kappa^S, \boldsymbol{\alpha}_{0:4}, \xi_{Amazon}, \xi_{Microsoft}, \sigma^E, \sigma^b, C_1^K, \sigma^K, \zeta]'$. I minimize a minimum distance objective function:

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} (\boldsymbol{\Psi}(\boldsymbol{\theta}) - \hat{\boldsymbol{\Psi}})' \mathbf{W} (\boldsymbol{\Psi}(\boldsymbol{\theta}) - \hat{\boldsymbol{\Psi}}) \tag{45}$$

The weighting matrix $\mathbf{W}$ is a diagonal weighting matrix in which the weights are chosen such that moments of different sizes are given roughly equal weight in the optimizer. In particular, moments above one receive a weight of 0.1, while all other moments receive a weight of 1. Observed statistics are averages over firm-years, in which sector level weights are given by the number of firms in each sector-year in the matched data, with the exception of entry costs, which are averaged across firms within each year.

Moment fit. I report targeted moments in Table C.2. Observed statistics are averaged over firm-years, with the standard deviation from 80 block-bootstrapped samples reported in parentheses. The sources of uncertainty come from sampling error in my sample of consumers. I do not resample suppliers; for supplier moments, uncertainty comes from sampling error in consumers' datacenter choices.

Table C.2: Empirical vs. simulated moments

Moment	Data	Simulation
Adoption rate 2018 ($\frac{\sum_f \cdot N_{2018} \mathbb{1}\{b_{ft} \neq \varnothing\}}{\sum_f N_{2018}}$)	0.1156	0.1044
	(0.0009)	
Adoption rate 2019 ($\frac{\sum_f \cdot N_{2019} \mathbb{1}\{b_{ft} \neq \varnothing\}}{\sum_f N_{2018}}$)	0.1335	0.1099
	(0.0007)	
Adoption rate 2020 ($\frac{\sum_f \cdot N_{2020} \mathbb{1}\{b_{ft} \neq \varnothing\}}{\sum_f N_{2018}}$)	0.2052	0.1336
	(0.0008)	

Continued on next page

Moment	Data	Simulation
Adoption rate 2021 ($\frac{\sum_f \cdot N_{2020} \mathbb{1}\{b_{ft} \neq \varnothing\}}{\sum_f N_{2018}}$)	0.2036	0.2024
	(0.0008)	
Bundle autocorrelation ($\mathbb{E}[\mathbb{1}\{b_{ft} = b_{ft-1}\} \vert b_{ft} \neq \varnothing]$)	0.9371	0.8463
	(0.0012)	
Multicloud rate ($\mathbb{E}[\vert \mathbb{1}\{\#j \ni b_{ft} > 1\} \vert b_{ft} \neq \varnothing]$)	0.3384	0.4166
	(0.0026)	
Exit rate ($\mathbb{E}[\mathbb{1}\{b_{ft} = \varnothing\} \vert b_{ft-1} \neq \varnothing]$)	0.0338	0.0547
	(0.0006)	
Size premium ($\mathbb{E}[K_{ft} \vert b_{ft} \neq \varnothing] / \mathbb{E}[K_{ft} \vert b_{ft} = \varnothing]$)	2.4961	1.6047
	(0.0401)	
Mean investment rate ($\mathbb{E}[\frac{I_{ft}}{K_{ft}}]$)	0.1576	0.2260
	(0.0006)	
Investment rate variance ($Var(\frac{I_{ft}}{K_{ft}})$)	0.0844	0.0526
	(0.0011)	
Own price correlation (new entrants) ($\mathbb{E}[p_{fbt} \cdot \mathbb{1}\{b_{ft}\} \vert b_{ft-1} = \varnothing]$)	0.0559	0.0696
	(0.0005)	
Own price corr. (no Microsoft server)		
($\mathbb{E}[p_{fbt} \cdot \mathbb{1}\{b_{ft}\} \vert b_{ft-1} = \varnothing, server_{ft} = 0]$)	0.0557	0.0623
	(0.0010)	
Own quality correlation (new entrants)		
($\mathbb{E}[\log(q_{bft}) \cdot \mathbb{1}\{b_{ft}\} \vert b_{ft-1} = \varnothing]$)	1.8008	1.8252
	(0.0020)	
Own quality corr. (no Microsoft server)		
($\mathbb{E}[\log(q_{bft}) \cdot \mathbb{1}\{b_{ft}\} \vert b_{ft-1} = \varnothing, server_{ft} = 0]$)	1.8245	1.8309
	(0.0049)	
Azure-Microsoft correlation		
($\mathbb{E}[\mathbb{1}\{Microsoft \in b_{ft}\} \cdot \mathbb{1}\{server_{ft}\} \vert b_{ft-1} = \varnothing]]$)	0.5997	0.5934
	(0.0030)	
Amazon entrant share (Microsoft server)		
($\mathbb{E}[\mathbb{1}\{Amazon \in b_{ft}\} \vert b_{ft-1} = \varnothing, server_{ft} = 0]$)	0.1923	0.1870
	(0.0018)	
Microsoft entrant share (Microsoft server)		
($\mathbb{E}[\mathbb{1}\{Microsoft \in b_{ft}\} \vert b_{ft-1} = \varnothing, server_{ft} = 0]$)	0.6793	0.7129
	(0.0024)	

Continued on next page

Moment	Data	Simulation
Amazon entrant share (no Microsoft server)		
$(\mathbb{E}[\mathbb{1}\{Amazon \in b_{ft}\}\|b_{ft-1} = \varnothing, server_{ft} = 0])$	0.3348	0.3461
	(0.0053)	
Microsoft entrant share (no Microsoft server)		
$(\mathbb{E}[\mathbb{1}\{Microsoft \in b_{ft}\}\|b_{ft-1} = \varnothing, server_{ft} = 0])$	0.4099	0.4519
	(0.0050)	
Multi-cloud size premium $(\mathbb{E}[K_{ft}\|\|b_{ft}\| > 1]/\mathbb{E}[K_{ft}\|\|b_{ft}\| = 1])$	1.4987	0.9423
	(0.0624)	
Amazon price $(\mathbb{E}[p_{Amazon,t}])$	0.0516	0.0603
	(0.0000)	
Microsoft price $(\mathbb{E}[p_{Microsoft,t}])$	0.0583	0.0848
	(0.0001)	
Google price $(\mathbb{E}[p_{Google,t}])$	0.0454	0.0768
	(0.0000)	
Entry rate $(\mathbb{E}[\mathbb{1}\{b_{ft} \neq \varnothing\}\|b_{ft-1} = \varnothing])$	0.0617	0.0117
	(0.0003)	

Note: Standard deviation of empirical moments calculated across 80 block bootstrapped samples, which reflect sampling uncertainty. Supplier price moments are computing using suppliers' optimization problem, while the remaining simulated moments come from consumer firms' optimization problem.

C.3 Computational details

In this appendix I provide details on the simulation for both the structural parameter estimation and counterfactual scenarios.[92] The estimation is complex, requiring thousands of nested fixed point iterations over firm value functions as well as guesses of the price paths. Even supplying analytical derivatives for suppliers and modeling simple parametric price forecasts, the optimizer often fails to converge far from the true parameter values. For example, suppliers' internalization of multihoming complementarities means that their pricing game is not always supermodular across the entire set of parameter values considered. To speed convergence, I smooth entry, maintenance, and investment costs using i.i.d. shocks. In addition, sectors in which $\beta_l + \beta_m \geqslant 0.95$ can produce instability for very large or very small values of (K_{ft}, Ω_{ft}). To speed convergence I multiply these by 0.9.[93]

[92] The author benefited from Claude Opus 4.0 through 4.8 and Claude Fable 5 to optimize code.

[93] This binds in services but not in either manufacturing sector. In contrast, Asker et al. (2014) forces returns to scale of the revenue production function to be 0.75, and Chen and Xu (2023) study an industry in which the capital share is large.

Despite these techniques, convergence failures remain. To remedy this, I follow Gortmaker (2024) and begin my point estimates using Bayesian optimization, a derivative-free global approach that can handle well noisy objectives. I use the `bayesopt` function in MATLAB, which constructs a surrogate Gaussian Process model from the evaluated objective and uses an acquisition function to select the next best parameter values, balancing exploitation and exploration. I use the expected improvement plus acquisition function. I find that an active set size of 100 and exploration ratio of 0.08 work well for my optimizer. I use the settings from Gortmaker (2024): 10 sets of 50 random initial evaluations within wide parameter bounds, followed by 150 Bayesian iterations for each sets. Failure rates reach up to 40% across initial guesses, though are far lower near the global optimum.[94]

After global optimization, I fine-tune with a round of local optimization starting from the parameters with the smallest objective across Bayesian evaluations, using the trust-based local refinement algorithm `fmincon` and the SQP algorithm. As an additional guard I apply a feasibility constraint that analyzes whether the MME loop has failed to converge:

$$\left| \frac{||\hat{M}_\psi - M_\psi||}{\epsilon^{tol}} \right| - 1 \leqslant 10^{-5}$$

Near the global optimum the constraint is never binding. For each bootstrap iteration, I begin with local optimization, warm starting each iteration at the point estimate and switching from central to forward differences in the SQP algorithm in the interest of tractability.

D Results appendix

In this appendix, I present additional results and robustness checks for my main specification. Appendix D.1 presents several alternative production function specifications. Appendix D.2 provides the long run productivity distributions for firms using 3 cloud suppliers, and Appendix D.3 presents cloud cost summary statistics for the services and high-tech manufacturing industries. Appendix D.4 compares model estimates to industry reports about the cost and benefits of IaaS adoption. Finally, Appendix E presents additional counterfactual results and robustness checks.

D.1 Additional production function estimates

This appendix presents a number of robustness checks to my baseline production function. Appendix D.1.1 estimates a more flexible productivity process for the gains from single- and

[94] If any parameter guess produces at least 10 supplier-years at which any supplier j's markup is $> €50$, I terminate the guess at outer iteration 40 and return an objective of `NaN`.

multihoming firms Appendix D.1.2 presents results for a translog production function, Appendix D.1.3 allows the productivity effects of cloud computing to vary over time to account for learning-by-doing, Appendix D.1.4 seeks to address measurement error in Aberdeen's measure of IT investment by using total capital investment as my first-stage proxy, and Appendix D.1.5 uses the full set of De Loecker et al. (2016) quality shifters.

D.1.1 Differential effects of single- and multihoming

In the baseline specification, I assume that the productivity gains from multihoming are a multiple of the number of providers used. In this appendix, I instead estiamte the following model

$$\omega_{ft} = \gamma_0 + \gamma_1 \omega_{ft-1} + +\gamma_2 \omega_{ft-1}^2 + +\gamma_3 \omega_{ft-1}^3 \tag{46}$$

$$+ \gamma_4 \mathbb{1}\{\#j = 1\} + \gamma_5 \mathbb{1}\{\#j > 1\} + \gamma_6 \cdot \mathbb{1}\{t = 2020\} + v_{ft} \tag{47}$$

Then $\gamma_5 - \gamma_4$ is the marginal gain of multihoming. Results are presented in Table D.1. In all sectors, the marginal return to mutlihoming is positive but displays a lower marginal return than single-homing. F-tests for each sector reject the null that the coefficients are identical.

Table D.1: Productivity evolution with flexible multihoming effect

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Productivity intercept (γ_0)	0.163 (0.022)***	0.155 (0.018)***	0.140 (0.003)***
Linear AR(1) term (γ_1)	0.714 (0.050)***	0.683 (0.045)***	0.498 (0.011)***
Quadratic AR(1) term (γ_2)	0.170 (0.037)***	0.230 (0.036)***	0.509 (0.013)***
Cubic AR(1) term (γ_3)	-0.034 (0.009)***	-0.057 (0.009)***	-0.167 (0.001)***
Single cloud adoption (γ_4)	0.007 (0.001)***	0.007 (0.001)***	0.004 (0.001)***
Multi-cloud adoption (γ_5)	0.011 (0.002)***	0.008 (0.002)***	0.005 (0.001)***
COVID-19 shock (γ_6)	-0.049 (0.001)***	-0.054 (0.001)***	-0.034 (0.001)***
Observations	14,893	37,948	153,230
Firms	4,877	12,791	70,321

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) wage control for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

D.1.2 Translog production function

My baseline model uses a Cobb-Douglas production function in order to allow me to partial out labor and materials from the firm's period profit function, dramatically simplifying the state space. To test the impact of a more flexible production function on my estimates, I estimate the following translog specification using the control function approach in Ackerberg et al. (2015):

$$y_{ft} = \omega_{ft} + \beta_l l_{ft} + \beta_k k_{ft} + \beta_m m_{ft} + \tag{48}$$

$$\beta_{ll} l_{ft}^2 + \beta_{kk} k_{ft}^2 + \beta_{mm} m_{ft}^2 + \tag{49}$$

$$\beta_{lk} l_{ft} k_{ft} + \beta_{km} k_{ft} m_{ft} + \beta_{lm} l_{ft} m_{ft} + \epsilon_{ft} \tag{50}$$

The moment conditions for the productivity process are identical. The moment conditions to identify $(\boldsymbol{\beta}, \boldsymbol{\vartheta})$ become:

$$\mathbb{E}[v_{ft}(\boldsymbol{\beta}, \boldsymbol{\vartheta}) \mathbf{I}_{ft}^{trans}] = 0 \tag{51}$$

where $\mathbf{I}_{ft}^{trans} = [k_{ft}, m_{ft-1}, l_{ft-1}, k_{ft}^2, m_{ft-1}^2, l_{ft-1}^2, m_{ft-1} l_{ft}, m_{ft} k_{ft}, l_{ft} k_{ft}, w_{ft-1}^l]'$. Production function result are given in Table D.2 and productivity estimates in Table D.3.

Table D.2: Production function: Translog

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Labor (β_l)	0.260 (0.166)	0.280 (0.208)	0.329 (0.285)
Materials (β_m)	0.708*** (0.182)	0.697*** (0.185)	0.701*** (0.258)
Capital (β_k)	0.023 (0.023)	0.023 (0.023)	0.014 (0.029)
Labor^2 (β_{ll})	-0.029 (0.112)	0.073 (0.279)	0.051 (0.061)
Materials^2 (β_{mm})	-0.033 (0.088)	0.067 (0.157)	0.043 (0.072)
Capital^2 (β_{kk})	0.001 (0.010)	0.002 (0.011)	0.012 (0.008)
Materials×labor (β_{ml})	0.0627† (0.036)	-0.136*** (0.041)	-0.096 (0.139)
Materials×capital (β_{mk})	0.003 (0.045)	-0.005 (0.054)	-0.002 (0.024)
Labor×capital (β_{lk})	-0.004 (0.036)	-0.001 (0.041)	-0.021 (0.041)
Wage (ϑ_w)	0.296 (0.470)	0.309 (0.352)	0.163 (0.367)
Observations	14,893	37,948	153,230
Firms	4,877	14,491	70,321

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) controls for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

Table D.3: Productivity evolution: Translog

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Productivity intercept (γ_0)	106.1 (9.661)***	0.157 (0.020)***	0.243 (0.007)***
Linear AR(1) term (γ_1)	83.82 (7.750)***	0.680 (0.048)***	0.424 (0.017)***
Quadratic AR(1) term (γ_2)	22.60 (2.071)***	0.234 (0.038)***	0.432 (0.014)***
Cubic AR(1) term (γ_3)	-2.008 (0.184)***	-0.059 (0.010)***	-0.110 (0.003)***
Cloud adoption (γ_4)	0.002 (0.0003)***	0.001 (0.0002)***	0.0009 (0.0001)***
COVID-19 shock (γ_5)	-0.052 (0.001)***	-0.053 (0.001)***	-0.028 (0.004)***
Observations	14,893	37,948	153,230
Firms	4,877	14,491	70,321

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) controls for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

D.1.3 Learning-by-doing and productivity

Next, I allow for the productivity impact of cloud computing adoption to vary over time, to model learning by doing effects (Brand et al. (2024), cf. Benkard (2000)) or unobserved intangible investments (Brynjolfsson et al. (2021)). My modified productivity process is as follows:

$$\omega_{ft} = \gamma_0 + \gamma_1 \omega_{ft-1} + +\gamma_2 \omega_{ft-1}^2 + +\gamma_3 \omega_{ft-1}^3 \tag{52}$$

$$+ \gamma_\tau \tau_{IaaS} + \gamma_5 \cdot \mathbb{1}\{t = 2020\} + v_{ft} \tag{53}$$

where τ_{IaaS} is the number of years a firm has used cloud computing.[95] Results are presented in Tables D.4 and D.5. In both manufacturing and services, the variable is positive and significant, indicating that the productivity effects of cloud computing are rising with additional use.

[95] Compare to the exporting variable in de Loecker (2013).

Table D.4: Production function: Learning by doing

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Labor (β_l)	0.129 (0.081)	0.307 (0.048)***	0.231 (0.036)***
Materials (β_m)	0.583 (0.031)***	0.477 (0.054)***	0.720 (0.184)***
Capital (β_k)	0.092 (0.028)**	0.079 (0.010)***	0.043 (0.008)***
Wage (ϑ_w)	0.148 (0.064)*	0.310 (0.046)***	0.224 (0.084)**
Observations	14,893	37,948	153,230
Firms	4,877	14,491	70,321

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) controls for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

Table D.5: Productivity evolution: Learning by doing

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Productivity intercept (γ_0)	0.163 (0.021)***	16.92 (3.864)***	2.01 (0.034)***
Linear AR(1) term (γ_1)	0.722 (0.047)***	-12.85 (4.016)**	-1.013 (0.053)***
Quadratic AR(1) term (γ_2)	0.159 (0.033)***	3.697 (1.391)**	0.682 (0.028)***
Cubic AR(1) term (γ_3)	-0.031 (0.008)***	-0.320 (0.160)*	-0.083 (0.005)***
Cloud adoption (γ_4)	0.002 (0.000)***	0.001 (0.000)***	0.002 (0.000)***
COVID-19 shock (γ_5)	-0.029 (0.001)***	-0.026 (0.001)***	-0.035 (0.001)***
Observations	14,893	37,948	153,230
Firms	4,877	14,491	70,321

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) controls for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

D.1.4 Total investment proxy variable

Aberdeen estimates the value of IT investments in euros based on its computation of the value of PCs, servers, and storage devices reported by firms in interviews. To account for potential measurement error (Tambe and Hitt (2012); Brynjolfsson and Hitt (2003)), I re-run my baseline regressions but use total capital investment as my proxy variable for the first stage, as is standard in many IO applications (Ackerberg et al. (2015); Olley and Pakes (1996)).[96] Results are presented in Tables D.6 and D.7. Even though median IT investment is 19.1%, 16.4%, and 18.9% of total investment in manufacturing, services, and high-tech

[96]However, He et al. (2026) compare Aberdeen's CITDB data on IT spending with regulatory disclosures on IT spending for US banks and finds a high degree of correlation.

manufacturing, respectively, differences from the baseline model are minimal.

Table D.6: Production function: Total investment proxy

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Labor (β_l)	0.191 (0.083)*	0.290 (0.047)***	0.343 (0.044)***
Materials (β_m)	0.608 (0.031)***	0.494 (0.052)***	0.686 (0.263)**
Capital (β_k)	0.064 (0.020)**	0.080 (0.010)***	0.012 (0.036)
Wage (ϑ_w)	0.214 (0.056)***	0.300 (0.044)***	0.048 (0.122)
Observations	15,903	46,631	164,500
Firms	4,974	14,995	73,127

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) controls for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

Table D.7: Productivity evolution: Total investment proxy

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Productivity intercept (γ_0)	0.177 (0.028)***	0.191 (0.018)***	-0.013 (0.001)***
Linear AR(1) term (γ_1)	0.661 (0.007)***	0.589 (0.045)***	-0.877 (0.002)***
Quadratic AR(1) term (γ_2)	0.223 (0.052)***	0.308 (0.035)***	0.160 (0.004)***
Cubic AR(1) term (γ_3)	-0.052 (0.013)***	-0.079 (0.001)***	-0.137 (0.004)***
Cloud adoption (γ_4)	0.004 (0.001)***	0.004 (0.001)***	0.002 (0.001)***
COVID-19 shock (γ_5)	-0.045 (0.001)***	-0.049 (0.001)***	-0.039 (0.001)***
Observations	15,903	46,631	164,500
Firms	4,974	14,995	73,127

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) controls for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

D.1.5 Including exporting and market share quality shifters

Existing work in industrial organization and trade finds significant productivity effects from firms choosing to export (de Loecker (2013)). In addition, De Loecker et al. (2016) find that firm-level market share to be correlated with output quality. To model this I re-estimate my model and modify the first stage of Ackerberg et al. (2015) as follows:

$$\tilde{\omega}_{ft}(\boldsymbol{\beta}, \boldsymbol{\vartheta}) = \hat{\phi}_{ft} - f(\mathbf{x}_{ft}; \boldsymbol{\beta}) - H(p^l_{ft}, ms_{\iota ft}, EXP_{ft}; \boldsymbol{\vartheta}) \tag{54}$$

where $ms_{\iota ft} \in [0,1]$ is firm f's market share in period t in industry ι measured at the 4-digit NACE code level and $EXP_{ft} \in \{0,1\}$ is an indicator for whether a firm f has positive export revenue in year t. Market shares are defined on the full Orbis sample[97], and export data comes from Orbis. The moment conditions for the productivity process are identical. The moment conditions to identify $(\boldsymbol{\beta}, \boldsymbol{\vartheta})$ become:

$$\mathbb{E}[v_{ft}(\boldsymbol{\beta}, \boldsymbol{\vartheta})\mathbf{I}_{ft}^{quality}] = 0 \tag{55}$$

where $\mathbf{I}_{ft}^{quality} = [k_{ft}, m_{ft-1}, l_{ft-1}, w_{ft-1}^{l}, EXP_{ft-1}, ms_{\iota ft-1}]'$. Results for the production function are presented in Table D.8 and for the productivity process in Table D.9.

Table D.8: Production function: Additional quality controls

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Labor (β_l)	0.246 (0.039)***	0.205 (0.031)***	0.162 (0.070)*
Materials (β_m)	0.629 (0.033)***	0.671 (0.051)***	0.559 (0.052)***
Capital (β_k)	0.039 (0.019)*	0.024 (0.018)	0.087 (0.028)**
Wage (ϑ_w)	0.212 (0.019)***	-0.090 (0.134)	0.247 (0.152)
Market share (ϑ_{ms})	0.268 (0.182)	0.471 (0.145)**	0.391 (0.152)*
Exporting (ϑ_{EXP})	1.367 (0.423)**	0.782 (0.389)*	1.470 (0.495)**
Observations	14,883	37,943	153,195
Firms	4,876	14,489	70,319

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) controls for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

[97] Missing revenue in Orbis is supplemented with Aberdeen data.

Table D.9: Productivity evolution: Additional quality controls

Variable	High-tech manufacturing	Low-tech manufacturing	Services
Productivity intercept (γ_0)	-0.011 (0.013)	-0.130 (0.003)***	0.215 (0.003)***
Linear AR(1) term (γ_1)	0.736 (0.020)***	0.733 (0.015)***	0.712 (0.009)***
Quadratic AR(1) term (γ_2)	-0.069 (0.029)*	-0.021 (0.035)	0.017 (0.012)
Cubic AR(1) term (γ_3)	-0.020 (0.036)***	-0.024 (0.024)	0.016 (0.004)
Cloud adoption (γ_4)	0.005 (0.006)	0.005 (0.002)†	0.003 (0.002)
COVID-19 shock (γ_5)	-0.046 (0.010)***	-0.042 (0.004)***	-0.038 (0.004)***
Observations	14,883	37,943	153,195
Firms	4,876	14,489	70,319

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\dagger}p < 0.1$

Note: Regressions are run using the method from Ackerberg et al. (2015) with the de Loecker (2013) correction for cloud adoption in productivity and De Loecker et al. (2016) controls for input quality. Standard errors are from 200 block bootstrap samples at the firm level.

D.2 Steady state productivity distribution for multihoming firms

Figure D.1 compares the steady state productivity distribution for each sector, comparing non-adopting firms with those that have adopted all 3 cloud products.

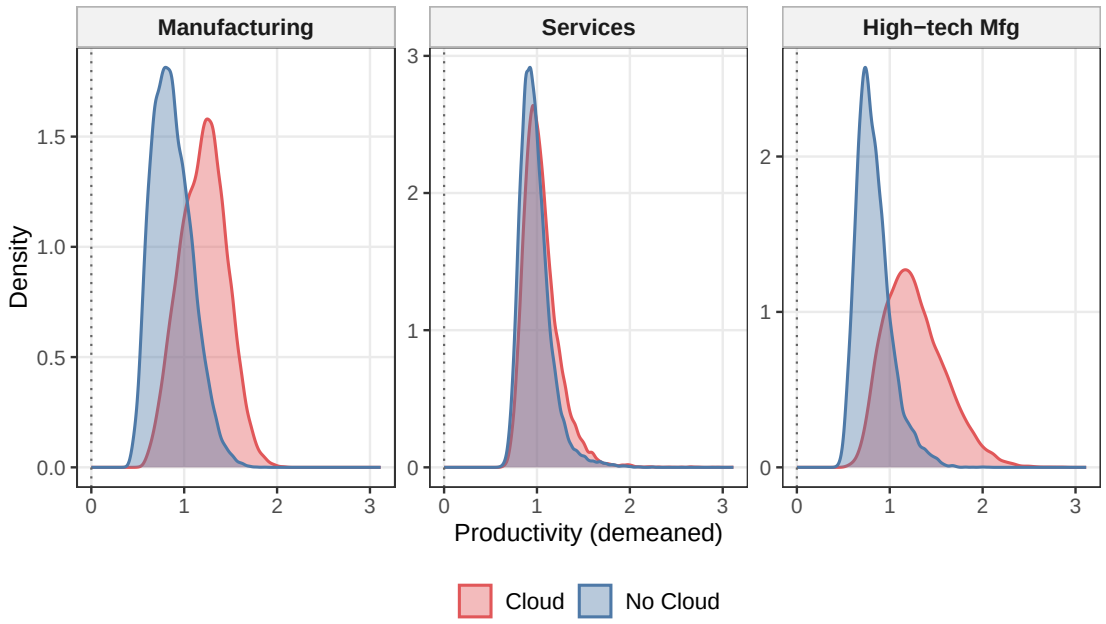


Figure D.1: Steady state productivity distribution (0 vs. 3 suppliers)

D.3 Cost distribution for remaining sectors

In this appendix, I present the estimates corresponding to Table 9 for the services and high-tech manufacturing sectors.

Table D.10: Services sector cloud cost distribution

Maintenance costs					
Supplier	Minimum	10th Percentile	Median	90th Percentile	Maximum
Microsoft	0.0616	0.0829	0.1697	0.3365	0.7759
Amazon	0.0960	0.1291	0.2643	0.5241	1.2086
Google	0.1335	0.1796	0.3676	0.7289	1.6810
Entry costs					
Year	Minimum	10th Percentile	Median	90th Percentile	Maximum
2017	1.6772	2.2567	4.6187	9.1577	21.1185
2021	0.1179	0.1586	0.3246	0.6436	1.4843

Note: Costs computed for firms used in the simulation, in millions of real euros. Supplier prices and qualities are computed using 2017 values.

Table D.11: High-tech manufacturing sector cloud cost distribution

Maintenance costs					
Supplier	Minimum	10th Percentile	Median	90th Percentile	Maximum
Microsoft	0.0472	0.1081	0.2573	0.8492	2.2580
Amazon	0.0735	0.1684	0.4008	1.3226	3.5171
Google	0.1022	0.2343	0.5574	1.8396	4.8919
Entry costs					
Year	Minimum	10th Percentile	Median	90th Percentile	Maximum
2017	1.2837	2.9429	7.0029	23.1112	61.4561
2021	0.0902	0.2068	0.4922	1.6244	4.3194

Note: Costs computed for firms used in the simulation, in millions of real euros. Supplier prices and qualities are computed using 2017 values.

D.4 Comparison to industry estimates of IaaS costs and benefits

In this appendix, I use the estimates from Forrester (North et al. (2019)) to perform a back of the envelope calculation of the real-world adoption costs and productivity gains from adopting IaaS and compare them to my estimates. The report conducted interviews with 6 high-ranking employees to create an estimate of cost and benefits of adopting Azure IaaS for 3 years for a representative firm with €3573 million in annual revenue. The report separately lists usage costs (fixed plus 3 years of maintenance) as well as gains from both lower IT costs and additional revenue.

For the Forrester data, I convert all values to euros using the mean 2019 exchange rate and deflate them to real euros. For my firm sample, for each sector i, I compute

$$median\{K_{ft}|Y_{ft} \in [0.9 \times 3573, 1.1 \times 3573]\}$$

and calculate the relevant costs and increases in profits for each sector, as well as a weighted average of my sample. For productivity, I take the median for each sector.[98] Results are displayed in Table D.12. While costs are lower in two sectors and profit increases are slightly larger, this may be due to the fact that I ignore reoptimization of investment, which would be on average larger for firms using cloud computing.

[98] Although cost functions are the same across sectors, data requirements at each level of output as well as productivity gains vary by sector.

Table D.12: Comparison of costs and benefits to Forrester estimates

	Forrester estimate	High-tech manufacturing	Low-tech manufacturing	Services
Costs	€4.88 M	€5.97 M	€0.44 M	€1.58 M
Profit change	0.51%	1.57%	1.80%	3.75%

Note: Costs computed for firms with capital stock $median\{K_{ft}|Y_{ft} \in [0.9 \times 3573, 1.1 \times 3573]\}$. Profit increases computed at median productivity, and I assume $server_{ft} = 1$. Profit changes do not account for re-optimization of capital investment. All costs and profits are computed using point estimates for 2019. Forrester profit estimates come from "Total Benefits" (North et al. (2019), p. 8) and costs from "Total Costs" (p. 17). All values in real 2015 euros.

E Additional counterfactual results

This appendix presents additional counterfactual results and robustness checks. Appendix E.1 presents a stylized one-period model to provide intuition for how the analyzed competition policies affect supplier pricing. Appendix E.2 provides the algorithm for solving for ζ^{CF}. Appendix E.3 provides additional robustness results for each counterfactual scenario. Finally, Appendix E.4 provides additional graphical results.

E.1 Stylized model of cloud adoption and pricing

In this appendix, I develop a simple stylized version of the model in Section 4 to build intuition about the impacts of competition policies on supplier pricing.

Consider a model in which there are two cloud suppliers, $j \in \{A, B\}$, who have already set egress fees $E_A(D_{ft}) = E_B(D_{ft}) = E$ and are now setting prices P_{jt}. There is a mass N_t of downstream firms of equal data stocks, normalized to 1, of which share n are new entrants, and m are multihoming firms. For simplicity, assume these are disjoint sets. Each supplier has a fraction $captive_j \in [0, 1]$ of the prior period single-homing consumers. Then, each supplier's captive set is given by

$$c_j = captive_j(N_t - n - m) \tag{56}$$

Finally, assume that a fraction of consumers use supplier A's on-premise servers ($server_{ft} = 1$) and receives a cost reduction $-\alpha_{server}, \alpha_{server} \geqslant 0$ for using A's cloud platform. Downstream firms have value functions

$$V_{ft} = \underbrace{\pi_{ft}}_{\text{profit}} - \underbrace{\alpha_{price} P_{jt}}_{\text{VM cost}} - \underbrace{E \cdot \mathbb{1}\{b_{ft} = j | b_{ft-1} = -j\}}_{\text{switching cost}} + \underbrace{\alpha_{server} \mathbb{1}\{j = A - \cap server_{ft} = 1\}}_{\text{server complementarity}} \tag{57}$$

$$- \underbrace{\lambda E \cdot \mathbb{1}\{b_{ft} = \{A, B\}\}}_{\text{multihoming egress cost}} + \epsilon_{ft} \tag{58}$$

where, as before, $\epsilon_{ft} \sim_{iid} Logistic(0, \sigma)$. For simplicty, set $\sigma = 1$. Defining $\tilde{u} \equiv \alpha_{price}(P_{Bt} - P_{At}) + \alpha_{server} S_t,$, where $S_t \equiv \mathbb{1}\{j = A - \cap server_{ft} = 1\}$, demand for firm A is given by

$$d_A = (n+m)\Lambda(\tilde{u}) + c_A \Lambda(\tilde{u} + E) + c_B \Lambda(\tilde{u} - E) \tag{59}$$

where $\Lambda(\cdot)$ is the logistic function and $d_B - N_t - d_A$. Normalizing costs to 0, profit for supplier A is given by

$$\pi_A^s = \underbrace{P_{At} d_A}_{\text{VM usage revenue}} + \underbrace{c_A E(1 - c_A \Lambda(\tilde{u} + E))}_{\text{Egress revenue from switchers}} + \underbrace{m\lambda E(n+m)\Lambda(\tilde{u})}_{\text{Egress revenue from multihomers}} \tag{60}$$

Taking the first-order conditions gives:

$$\frac{\partial \pi_{At}^s}{\partial P_{At}} = d_A - \alpha_{price} P_{At} \cdot \left((n+m)\tilde{\Lambda}(\tilde{u}) + c_A \tilde{\Lambda}(\tilde{u} + E) + c_B \tilde{\Lambda}(\tilde{u} - E) \right) \tag{61}$$

$$+ c_A \tilde{\Lambda}(\tilde{u} + E) - m\lambda E(n+m)\tilde{\Lambda}(\tilde{u}) \tag{62}$$

$$\text{where } \tilde{\Lambda}(\cdot) \equiv \Lambda(\cdot)(1 - \Lambda(\cdot)) \tag{63}$$

E.1.1 Effect of egress fees

What is the effect of lowering egress fees? To analyze this, consider the case where $\alpha_{server} = 0$, and assume $c_A = c_B$ such that suppliers are symmetric. Then, $\tilde{u} = 0$, so

$$\Lambda(\tilde{u}) = 1/2$$
$$\Lambda(\tilde{u} - E) = \Lambda(\tilde{u} + E) = \tfrac{1}{4} sech^2(\tfrac{E}{2})$$

Plugging in and solving, we have the symmetric equilibrium price:

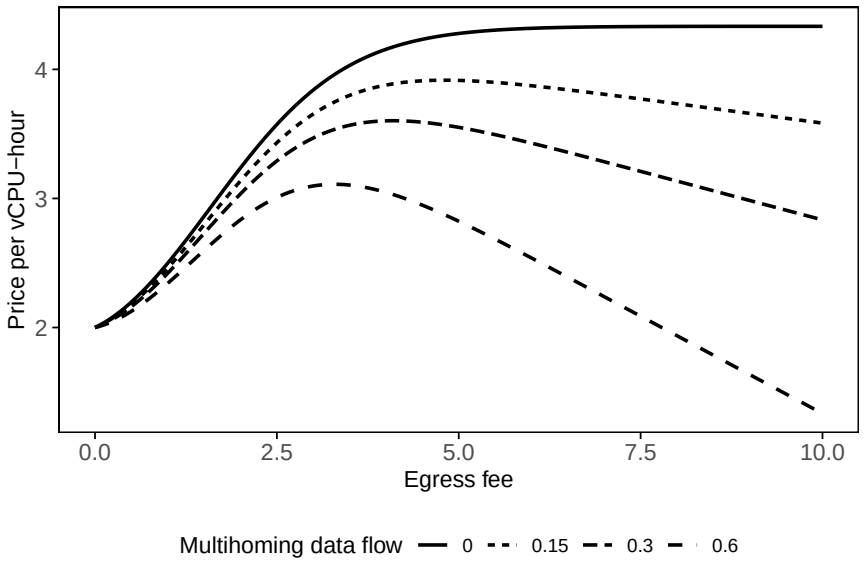


Figure E.1: Equilibrium price as egress fee varies in stylized model

Simulated price per vCPU-hour in stylized model under various levels of egress fee (E) and fraction of data that multihoming firms move between suppliers (λ). Parameters for simulation: $\alpha_{price} = 1, \alpha_{server} = 0, N_t = 1.3, m = 0.3, n = 0.3$.

$$P_{At} = P_{Bt} = \frac{2n + E[\frac{1}{2}(1-m)\frac{1}{4}sech^2(\frac{E}{2}) - \lambda m]}{\alpha_{price}[n + m + (1-m)\frac{1}{4}sech^2(\frac{E}{2})]} \tag{64}$$

Noting that $\frac{\partial sech^2(E/2)}{\partial E} \geqslant 0$ for all $E \geqslant 0$, we have three effects. If the share of multihoming firms is large enough, then prices fall with egress fees as these consumers become more valuable. For the captive firms who pay egress fees upon switching, there are two effects. At low E, higher egress fees result in a higher price (harvesting). At high enough E, no consumers are willing to switch supplier, such that the multihoming effect dominates and prices begin to fall. For a range of λ, I illustrate this graphically in E.1.

E.1.2 Effect of improved server integration

Next, consider the case in which $\alpha_{server} \neq 0$. In this case, as α_{server} rises, d_A rises relative to d_B, allowing A to charger a higher price and forcing B to charge a lower price. I illustrate this graphically in E.2.

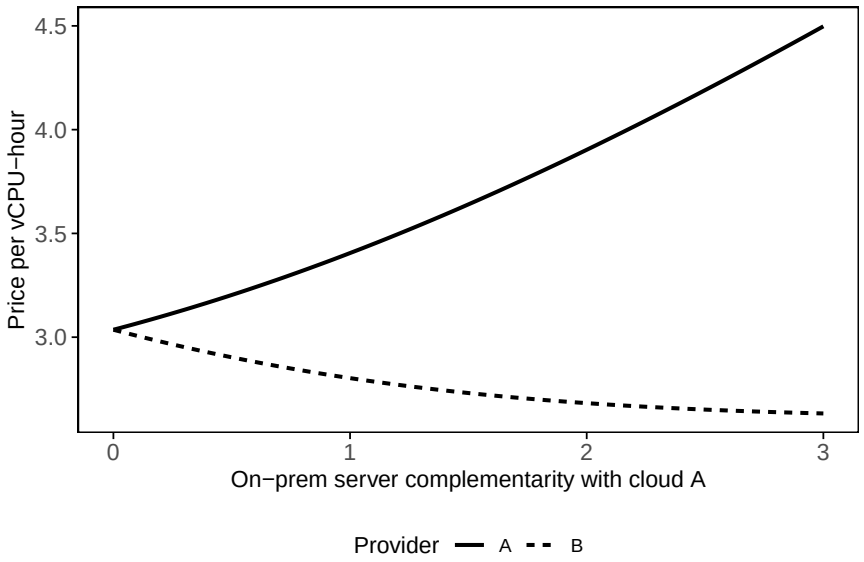


Figure E.2: Equilibrium price as server complementarity varies in stylized model

Simulated price per vCPU-hour in stylized model under various levels of server complementarity ($\sigma\alpha_{server}$). I assume that all firms have $S_t = 1$. Parameters for simulation: $\alpha_{price} = 1, \alpha_{server} = 0, N_t = 1.3, m = 0.3, n = 0.3$.

E.2 Deriving the counterfactual invest-harvest parameter

In this section, I describe the algorithm used to compute the counterfactual invest/harvest parameter, ζ^{CF}. I begin the algorithm at the estimated version of $\hat{\zeta}$ in order to obtain a counterfactual equilibrium near the empirical one.

1. Guess $\zeta^0 = \hat{\zeta}$.
2. Solve the model and set $\zeta^1 = mean[1 + \beta \frac{\partial \mathbb{E}_t[W_{jt+1}(s_{t+1})]/\partial P_{jt}}{\tilde{Q}_{jt}}]$ across $t \in [2017, ..., 2021]$.
3. For each iteration $k = 2, 3, \ldots$:
   (a) Solve the model under the counterfactual scenario and obtain $mean\left[\frac{\partial \mathbb{E}_t[W_{jt+1}(s_{t+1})]/\partial P_{jt}}{\tilde{Q}_{jt}}\right]$ across $t \in [2017, ..., 2021]$.
   (b) Set $\zeta^{k+1} = \frac{\zeta^k - \Gamma(\zeta^{k-1})(\zeta^k - \zeta^{k-1})}{\Gamma(\zeta^k) - \Gamma(\zeta^{k-1})}$, where $\Gamma(\zeta^k) = \zeta^k - 1 - (\zeta^{k-1} - 1) \cdot \frac{\frac{\partial \mathbb{E}_t[W_{jt+1}(s_{t+1})]/\partial P_{jt}}{\tilde{Q}_{jt}}|_{\zeta^k}}{\frac{\partial \mathbb{E}_t[W_{jt+1}(s_{t+1})]/\partial P_{jt}}{\tilde{Q}_{jt}}|_{\zeta^{k-1}}}$
4. If $|\zeta^{k+1} - \zeta^k| < \epsilon^\zeta = 1e-4$, terminate.
5. Proceed to iteration $k + 1$.

	ΔCS	ΔPS	ΔWelfare	ΔPrice	ΔAdoption	ΔMultihoming
Fixed ζ						
New entrant	5289.3	2960.3	8249.6	-0.018	1.33pp	5.89pp
	(1.84%)	(22.32%)	(2.77%)	(-24.57%)	(3.19%)	(19.56%)
Google-Microsoft merger	-5614.6	-777.3	-6391.8	0.042	-2.75pp	1.46pp
	(-1.99%)	(-5.85%)	(-2.16%)	(56.08%)	(-6.23%)	(5.39%)

Table E.1: Welfare effects of varying number of suppliers

Notes: Welfare changes relative to the estimated baseline over 2017-2031, in millions of real 2015 euros, averaged over 80 block bootstrap samples. Adoption the average annual share of firms using cloud computing. Multihoming is the average annual percentage of firms using multiple cloud suppliers. Price is the unweighted average annual price among suppliers $j \in \mathcal{J}$.

	ΔCS	ΔPS	ΔWelfare	ΔPrice	ΔAdoption	ΔMultihoming
Fixed ζ						
Price response	311.6	-122.5	189.2	-0.001	0.20pp	0.49pp
	(0.11%)	(-0.93%)	(0.06%)	(-1.28%)	(0.35%)	(1.56%)
No price response	104.0	-153.2	-49.2	–	0.22pp	-0.2800
	(0.04%)	(-1.15%)	(-0.02%)	(–)	(0.53%)	(-0.65%)

Table E.2: Welfare effects of egress fee ban

Notes: Welfare changes relative to the estimated baseline over 2017-2031, in millions of real 2015 euros, averaged over 80 block bootstrap samples. Adoption the average annual share of firms using cloud computing. Multihoming is the average annual percentage of firms using multiple cloud suppliers. Price is the unweighted average annual price among suppliers $j \in \mathcal{J}$.

E.3 Additional counterfactual results

In this appendix, I provide additional results for each counterfactual scenario. For the entrant and merger scenarios, I present results with ζ held fixed at its baseline (E.1). While the results for the entrant scenario are quite similar to tbe baseline, the merger scenario sees a far high price increase and welfare losses that are nearly an order of magnitude larger. In the merger scenario, ζ falls relative to the baseline, indicating that firms, especially the merged entity, are more interested in investing in a consumer base. This serves to lower prices compared to the fixed ζ scenario, ameliorating welfare losses.

For the egress fee scenario, I present results for holding prices fixed, as observed since the partial elimination of egress fees in 2024, and for holding ζ fixed the factual value (E.2). Holding prices fixed, we see that the industry experiences a small welfare loss - thus, the egress fee ban is welfare-improving only if it induces additional price competition among suppliers.

For the increased server interoperability scenario, I present results for varying $\varepsilon \in \{0.1, 0.2, 0.5\}$

	ΔCS	ΔPS	ΔWelfare	ΔPrice	ΔAdoption	ΔMultihoming
A: Updating ζ						
$\varepsilon = 0.1$	5.5	-5.7	-0.3	0.000	0.00pp	-0.05pp
	(0.00%)	(-0.04%)	(-0.00%)	(0.16%)	(0.00%)	(-0.09%)
$\varepsilon = 0.2$	105.1	1.2	106.3	0.000	0.04pp	-0.09pp
	(0.03%)	(0.01%)	(0.03%)	(-0.23%)	(0.11%)	(-0.19%)
B: Fixed ζ						
$\varepsilon = 0.1$	132.6	12.4	144.9	-0.000	0.06pp	-0.05pp
	(0.05%)	(0.09%)	(0.05%)	(-0.57%)	(0.14%)	(-0.11%)
$\varepsilon = 0.2$	296.0	28.8	324.8	-0.001	0.15pp	-0.11pp
	(0.10%)	(0.22%)	(0.11%)	(-1.29%)	(0.31%)	(-0.25%)
$\varepsilon = 0.5$	1140.6	121.6	1262.2	-0.004	0.57pp	-0.33pp
	(0.40%)	(0.92%)	(0.42%)	(-5.25%)	(1.21%)	(-0.88%)

Table E.3: Welfare effects of increased server interoperability

Notes: Welfare changes relative to the estimated baseline over 2017-2031, in millions of real 2015 euros, averaged over 80 block bootstrap samples. Adoption the average annual share of firms using cloud computing. Multihoming is the average annual percentage of firms using multiple cloud suppliers. Price is the unweighted average annual price among suppliers $j \in \mathcal{J}$.

and for holding ζ fixed at the factual value (E.3). The results indicate that welfare impacts of increased complementarity convex in how much complementarity is allowed - doubling ε more than doubles annual welfare gains.

E.4 Additional counterfactual figures

This appendix displays figures for additional counterfactual scenarios. Figure E.3 displays the evolution of consumer surplus from 2017 to 2031. Figure E.5 displays counterfactual price and adoption changes holding ζ fixed at the baseline estimate, and Figure E.5 displays annual changes in comsumer welfare under all four counterfactuals when holding ζ fixed.

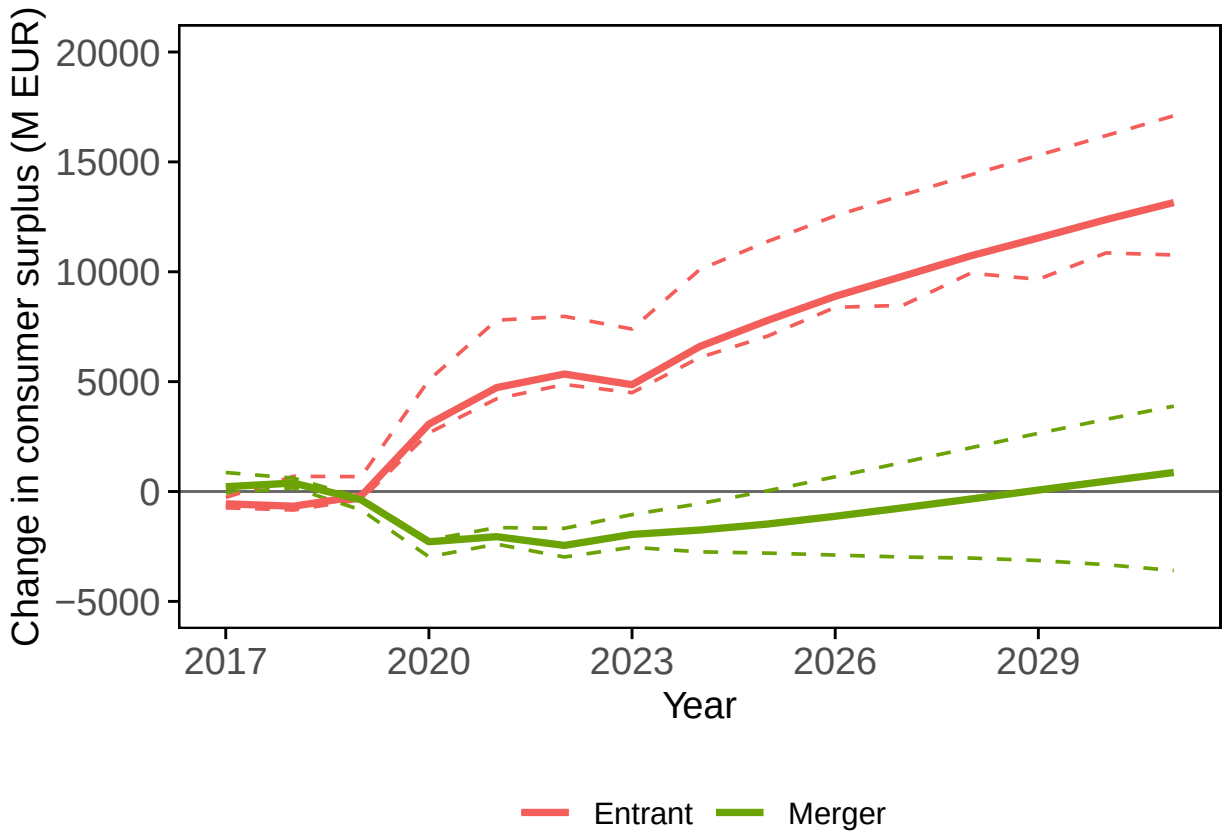
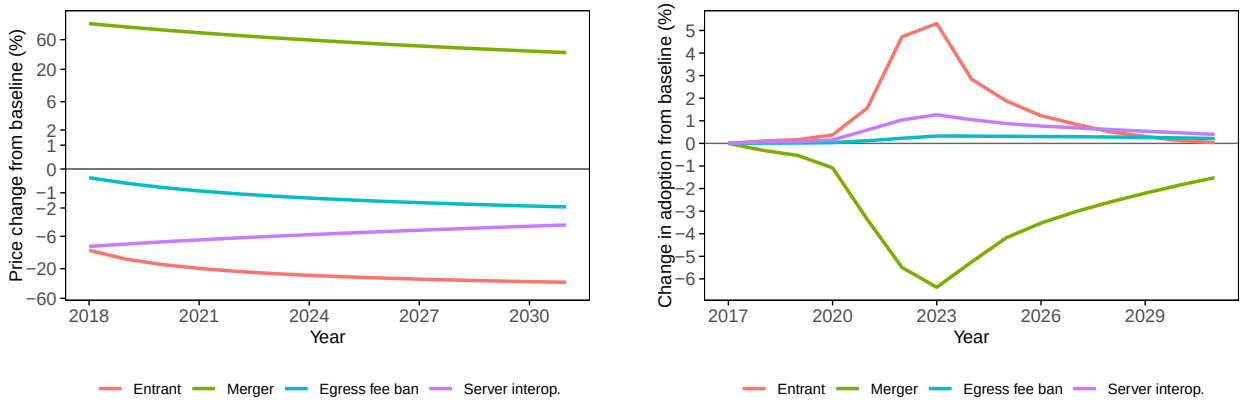


Figure E.3: Annual change in consumer surplus under changes in market structure

Annual change in consumer surplus, 2017-2031, with 95% CI from 80 block bootstrap samples. Values are in real 2015 euros.

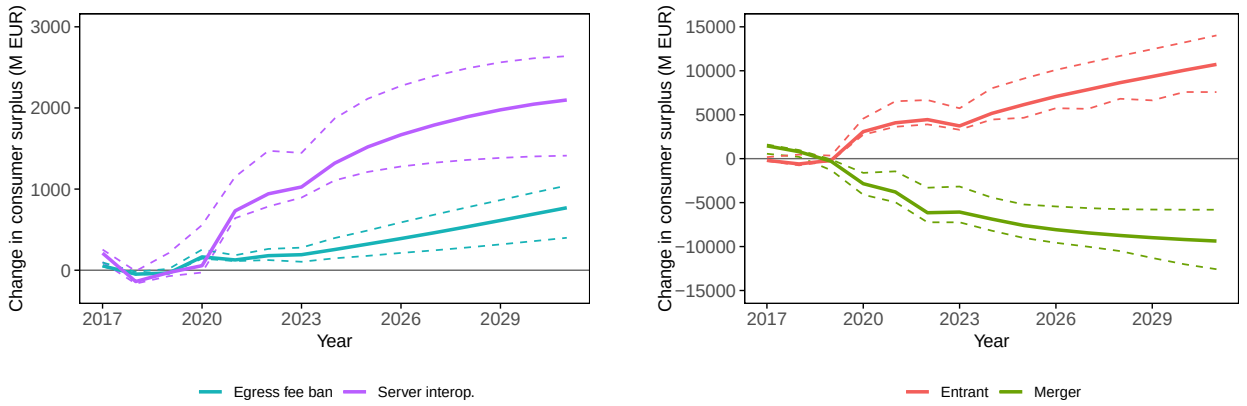


(a) Counterfactual price changes

(b) Counterfactual adoption changes

Figure E.4: Counterfactual dynamics holding ζ fixed: prices and adoption rates

Panel (a) plots the average annual change in price among all suppliers compared to the baseline, including the entrant in the relevant scenario. Panel (b) plots the average annual change in the fraction of firms using cloud computing compared to the baseline. Note that the y-axis is on a log scale. Annual values in (a) and (b) are the mean across 80 block bootstrap samples.



(a) Consumer surplus changes from competition policies

(b) Consumer surplus changes from changes in market structure

Figure E.5: Annual change in consumer surplus holding ζ fixed

Annual change in consumer surplus, 2017-2031, with 95% CI from 80 block bootstrap samples. Values are in real 2015 euros. Panel (a) plots the consumer surplus changes under the egress fee ban and increased Microsoft server interoperability ($\varepsilon = 0.5$) scenarios, and panel (b) plots the changes under the entrant and merger scenarios.

F Further institutional details

This appendix provides further description of the cloud computing services market in France. I first provide details on the IaaS market (Appendix F.1). Then, I provide further details on relevant competition policies for each of my counterfactuals (Appendix F.2).

F.1 The Infrastructure-as-a-Service market

Cloud providers offer three types of cloud services: software-as-a-service (SaaS), infrastructure-as-a-service (IaaS), and platform-as-a-service (PaaS). SaaS is the broadest term and encompasses various software leasing, including including Google Docs, Microsoft Office 365, and Dropbox. While many of these services are IaaS bundled with a software interface, their contracts are generally highly varied and encompass bundling rental of IT hardware with software licenses.[99] PaaS also abstracts away from the underlying hardware by focusing on an application layer such as serverless computing or containers. I ignore this segment in this study for two reasons. First, I do not observe the specific bundling of software and hardware that firms in my sample purchase, limiting my ability to study pricing. In addition, evidence suggests that IaaS is more common among consumers.[100]

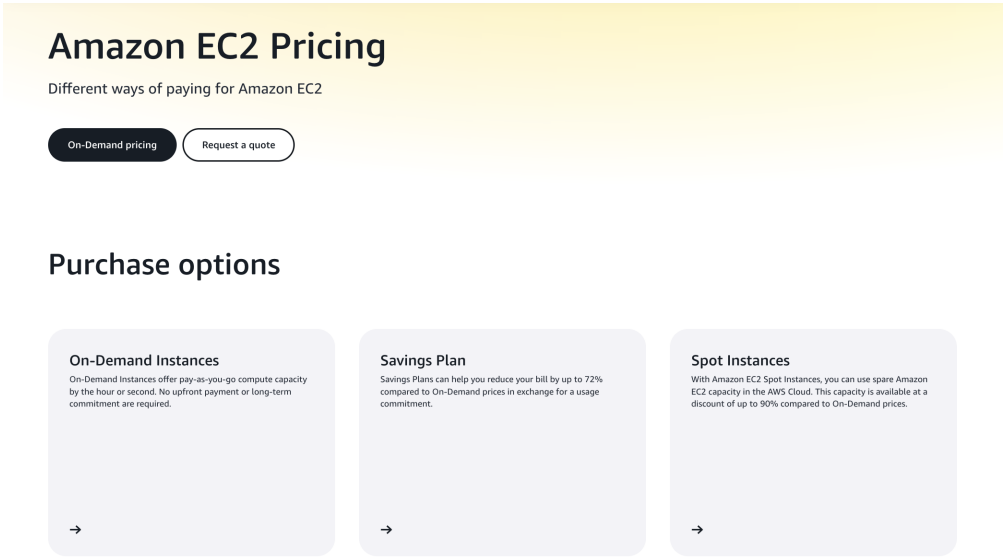
Finally, IaaS is the leasing of storage (usually as "blobs") or computation in the form of virtual machines (VMs). Users have a wide array of options for these - for the former, this includes data speeds, and for the latter this includes such factors as CPU clock speed, number of cores, and memory size. I focus on IaaS for two reasons. First, the economic mechanisms are clear - IaaS has the potential to lower the cost of computation through scale economies and allows firms to dynamically respond to productivity shocks, lowering IT capital adjustment frictions (Asker et al. (2014)). Second, the set of pricing options are fairly simple. For computation, firms can purchase on-demand instances, which are charged by the minute, or reserve computation for a period of 1 or 3 years for a reduced per unit price. Finally, firms can also purchase additional compute on the spot market, with the stipulation that the VM will be shut down if the firm is outbid for it.[101] In the data, most firm computation is done via on-demand prices, with a substantial minority of firms using reserved instances. A 2019 Flexera report found that 47% of AWS, 23% of Azure users, and 10% of Google Cloud users utilized reserved instances (Flexera (2019)). For storage,

[99] For example, Dropbox used Amazon's services until 2020, when it switched to purchasing its own data centers.

[100] See Brand et al. (2024) Appendix A. According to the CMA, in 2022, Azure's top product was compute, or IaaS virtual machines (Competition and Markets Authority (2025a), 17).

[101] Though some larger firms such as Netflix sign contracts with individual cloud providers, discussions with a social scientist working with a large cloud provider suggest such contracts are rare.

firms can purchase standard storage, with marginal price declining by volume, or inexpensive long-term storage that costs an additional amount for retrieval. The relative simplicity of IaaS pricing allows me to estimate firms' price elasticity of demand for cloud inputs, even though I do not observe firm-level prices in my data.



(a) Pricing options for VMs

Compute Savings Plans for Amazon EC2

Compute Savings Plans apply to EC2 Instance usage regardless of instance family, size, AZ, AWS Region, OS or tenancy.

Select a location type and region

Location Type: AWS Region — Region: Europe (Paris)

Select terms for your Compute Savings Plans

Term length: 1 year — Payment options: All Upfront

Select an operating system and tenancy to view rates

Operating system: Linux — Tenancy: Shared

Viewing 448 available instances

1 2 **3** 4 5 6 7 ... 23

Instance name	Savings Plans rate	Savings over On-Demand	On-Demand rate	Region	Operating system	Tenancy
m7gd.2xlarge	$0.3755	25%	$0.499	Europe (Paris)	Linux	Shared
m7gd.4xlarge	$0.7509	25%	$0.9979	Europe (Paris)	Linux	Shared
m7gd.8xlarge	$1.5018	25%	$1.9958	Europe (Paris)	Linux	Shared
m7gd.12xlarge	$2.2527	25%	$2.9938	Europe (Paris)	Linux	Shared
m7gd.16xlarge	$3.0036	25%	$3.9917	Europe (Paris)	Linux	Shared
m7gd.metal	$3.0036	25%	$3.9917	Europe (Paris)	Linux	Shared
m7i.large	$0.08324	29%	$0.1176	Europe (Paris)	Linux	Shared

(b) A sample of VM prices

Figure F.1: Pricing options for Amazon Web Services

F.2 Competition policy for cloud computing services

This appendix provides an overview of the competition policies relevant to cloud computing, focusing on areas of overlap across regulatory bodies.

Facilitating switching. A primary concern of multiple regulatory bodies are the barriers to switching and multihoming. The EU Data Act, mandates a graduate phaseout of egress fees: "From 11 January 2024 to 12 January 2027, providers of data processing services may impose reduced switching charges on the customer for the switching process" and then "[f]rom 12 January 2027, providers of data processing services shall not impose any switching charges on the customer for the switching process" (European Parliament and Council of the European Union (2023)). France's 2024 SREN law moves this process forward, capping egress fees to the marginal cost of data transfers in January 2026, which ARCEP has determined to be €0 (Arcep (2025)).

Other regulatory agencies have not yet banned egress fees. THe UK CMA finds that "technical and commercial barriers [to switching and multi-cloud] are features that give rise to an AEC [adverse effects on competition] in cloud services in the UK. These barriers lock customers into their initial choice of provider, which may not reflect their evolving needs and limits their ability to exercise choice of cloud provider. These barriers can restrict customers from responding to attractive offers or accessing innovative new services from another provider, leading to weaker competition between providers" (Competition and Markets Authority (2025b)). As a result, the CMA recommends commencing investigations into designating Microsoft and Amazon with strategic market status, granting the CMA statutory authority to impose tailored conduct requirements. Finally, both the Japan Fair Trade Commission (JFTC) and Korean Fair Trade Commission have concluded that egress fees are anticompetitive, though neither country has moved to restrict them (Japan Fair Trade Commission (2022); Korea Fair Trade Commission (2022)).

Ensuring software interoperability. The other major set of competition policies pursued by regulators concerns software interoperability. In addition to interoperability between cloud providers to facilitate switching and multihoming, regulators are also seeking to ensure interoperability between some provider j's software suite when run on the cloud of provider j' in order to prevent dominant firms from strategically bundling cloud computing with software products in order to foreclose rivals. The EU Data Act mandates that all cloud providers must "take all reasonable measures in their power to facilitate that the customer, after switching to a service covering the same service type, achieves functional equivalence in the use of the destination data processing service" (Competition and Markets Authority (2025b)). For example, when exiting a provider, "the provider of data processing services

shall, at the request of the customer, export all exportable data in a structured, commonly used and machine-readable format" (Competition and Markets Authority (2025b)).

The second form of interoperability concerns that between Microsoft server products and rival cloud platforms. The UK CMA finds that Microsoft's licencing practices raise the costs for consumers of using their server products on rival cloud platforms compared to Azure. Due to Microsoft's dominance in the server operating system market, with a 2018-19 market share of 72% (Amazon Web Services (2024)), the UK CMA finds that "Microsoft's licensing practices have the effect of reducing competition in cloud services markets by adversely impacting the competitiveness of AWS and Google in the supply of cloud services, particularly in competing for customers that purchase cloud services which use the relevant Microsoft software as an input" (Competition and Markets Authority (2025b), p. 16). Likewise, in 2026, the JFTC opened an investigation into Microsoft's suspected violation of the Anti-monopoly Act: namely, the firm is suspected of "(1) not allowing [consumers] to use the [server] Services in combination with Competing Cloud Services; or (2) modifying or setting trade terms of the Services so that costs for the users, etc. are increased when the Services are used in combination with Competing Cloud Services compared to the case in which the [server] Services are used on Azure" (Japan Fair Trade Commission (2026)).

References

Aberdeen Group (2017). *Aberdeen Technology Data Cloud: Data Dictionary*. San Diego, CA and Les Ulis, France, revision 3.14 edition. © 2017 Aberdeen Group. All rights reserved.

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Amazon Web Services (2024). Aws response to the cma's licensing working paper. Response to working paper, Competition and Markets Authority.

Amazon Web Services (n.d.). Rules for specifying cpu options for an amazon ec2 instance. `https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-cpu-options-rules.html`. Accessed: 2026-04-17.

Arcep (2025). Regulation of cloud computing services: Arcep adopts a recommendation on the interoperability and portability of cloud services. Press release. Press release 71-25.

Asker, J., Collard-Wexler, A., and De Loecker, J. (2014). Dynamic inputs and resource (mis) allocation. *Journal of Political Economy*, 122(5):1013–1063.

Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review*, 90(4):1034–1054.

Benkard, C. L., Jeziorski, P., and Weintraub, G. (2024). Transitional market dynamics in complex environments. NBER Working Paper 33045, National Bureau of Economic Research.

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

Boudreau, B. (2022). Price erosion remains the universal norm.

Brand, J. M., Demirer, M., Finucane, C., and Kreps, A. A. (2024). Firm productivity and learning in the digital economy: Evidence from cloud computing. Technical report, National Bureau of Economic Research.

Brynjolfsson, E. and Hitt, L. M. (2003). Computing productivity: Firm-level evidence. *Review of economics and statistics*, 85(4):793–808.

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.

Chen, Y. and Xu, D. Y. (2023). A structural empirical model of r&d, firm heterogeneity, and industry evolution. *The Journal of Industrial Economics*, 71(2):323–353.

Competition and Markets Authority (2025a). Appendix a: Demand for cloud services and how customers purchase cloud services. Appendix to final report, cloud services market investigation, UK Competition and Markets Authority. Appendix A to the CMA's final decision on the market investigation into public cloud infrastructure services.

Competition and Markets Authority (2025b). Cloud infrastructure services: Final decision report. Final decision report, Competition and Markets Authority.

Cortez, E., Bonde, A., Muzio, A., Russinovich, M., Fontoura, M., and Bianchini, R. (2017). Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 153–167.

Datacenters.com (2025). Search data centers: Colocation, bare metal, cloud. Online database / marketplace. Accessed: 2025-11-05.

de Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics*, 5(3):1–21.

De Loecker, J., Goldberg, P. K., Khandelwal, A. K., and Pavcnik, N. (2016). Prices, markups, and trade reform. *Econometrica*, 84(2):445–510.

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.

EC2Instances.info Contributors (2025). Ec2instances.info: Easy amazon ec2 instance comparison. Accessed: 2025-09-15.

European Parliament and Council of the European Union (2023). Regulation (eu) 2023/2854 of the european parliament and of the council of 13 december 2023 on harmonised rules on fair access to and use of data and amending regulation (eu) 2017/2394 and directive (eu) 2020/1828 (data act). Official Journal of the European Union, L 2023/2854. Published 22 December 2023; CELEX: 32023R2854.

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

Flexera (2019). Rightscale 2019 state of the cloud report from flexera. Technical report, Flexera. Survey of 786 technical professionals; licensed under Creative Commons Attribution4.0 International License.

Goettler, R. L. and Gordon, B. R. (2011). Does amd spur intel to innovate more? *Journal of Political Economy*, 119(6):1141–1200.

Gortmaker, J. (2024). Open source software policy in industry equilibrium. Technical report, Working paper, Tech. Rep.

Gowrisankaran, G. and Rysman, M. (2012). Dynamics of consumer demand for new durable goods. *Journal of political Economy*, 120(6):1173–1219.

Gowrisankaran, G. and Schmidt-Dengler, P. (2024). A computable dynamic oligopoly model of capacity investment. Technical report, National Bureau of Economic Research.

Hazell, J., Patterson, C., Sarsons, H., and Taska, B. (2022). National wage setting. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2022-150).

He, Z., Jiang, S., Xu, D., and Yin, X. (2026). Investing in lending technology: It spending in banking. *Management Science*.

Ifrach, B. and Weintraub, G. Y. (2017). A framework for dynamic oligopoly in concentrated industries. *The Review of Economic Studies*, 84(3):1106–1150.

Japan Fair Trade Commission (2022). Report on trade practices in cloud services sector. Report on fact-finding survey on trade practices by digital platform operators, Japan Fair Trade Commission. Report on Fact-Finding Survey on Trade Practices by Digital Platform Operators.

Japan Fair Trade Commission (2026). Jftc opens an investigation and seeks information and comments from third parties concerning the suspected violation of the antimonopoly act by microsoft corporation, etc. Press release.

Korea Fair Trade Commission (2022). Kftc announces results of cloud service market study. `https://www.ftc.go.kr/eng/downloadBbsFile.do?atchmnflNo=17414`. Press release, December 28, 2022 (hereinafter "S. Kor. FTC Cloud Study").

Lashkari, D., Bauer, A., and Boussard, J. (2024). Information technology and returns to scale. *American Economic Review*, 114(6):1769–1815.

Lee, S., Hwang, J., and Lee, K. (2022). Spotlake: Diverse spot instance dataset archive service. In *2022 IEEE International Symposium on Workload Characterization (IISWC)*, pages 242–255. IEEE.

Mach, W. and Schikuta, E. (2011). A consumer-provider cloud cost model considering variable cost. In *2011 IEEE Ninth International Conference on Dependable, Autonomic and Secure Computing*, pages 628–635. IEEE.

McCallum, J. C. and of Labor Statistics, U. B. (2024). Disk. Based on: John C. McCallum (2023), *Price and Performance Changes of Computer Technology with Time*; U.S. Bureau of Labor Statistics (2024), *US Consumer Prices*; with minor processing by Our World in Data.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society*, pages 995–1026.

Microsoft Azure (2019). Azure 2019 public dataset v2 – trace analysis. `https://github.com/Azure/AzurePublicDataset/blob/master/analysis/Azure%202019%20Public%20Dataset%20V2%20-%20Trace%20Analysis.ipynb`. Jupyter Notebook.

North, M. A., Cavallaro, R., and Davison, D. (2019). The total economic impact™ of microsoft azure iaas: Cost savings and business benefits enabled by azure infrastructure-as-a-service (iaas). Technical report, Forrester Consulting, commissioned by Microsoft. Forrester Research, Inc.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.

Pakes, A. and Pollard, D. (1989). Simulation and the asymptotics of optimization estimators. *Econometrica: Journal of the Econometric Society*, pages 1027–1057.

PassMark Software (2026a). Cpu benchmarks. `https://www.cpubenchmark.net/`. Accessed: 2026-04-17.

PassMark Software (2026b). Intel xeon e5-2660 @ 2.20ghz cpu benchmark. `https://www.cpubenchmark.net/cpu.php?cpu=Intel+Xeon+E5-2660+%40+2.20GHz&id=1219`. Accessed: 2026-04-17.

Tambe, P. and Hitt, L. M. (2012). The productivity of information technology investments: New evidence from it labor data. *Information systems research*, 23(3-part-1):599–617.

Yun, Y. (2026). Navigating the commons. Working paper, UCLA Anderson School of Management.