CESifo WORKING PAPERS

10546
2023

July 2023

E-Globalization and Trade Agreements

Phillip McCalman

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 10546*

E-Globalization and Trade Agreements

Abstract

The global success of online search engines and social media is due to their free access and high level of quality. However, these features are supported by a business model that exploits personal user data to provide targeted advertising services to third parties. Does this business model deliver socially desirable outcomes at the global and/or national level? To explore these questions, we characterize how a global monopoly platform chooses the level of privacy protection and service quality. When a platform operates a free service model it over-exploits personal information and underprovides quality compared to a global planner. Despite distortions along two dimensions, global welfare can be improved by a policy of enhanced privacy protection alone. In fact, it is likely that enhanced privacy protection will also induce higher platform quality. Furthermore, when privacy policies are set at the national level, large countries tend to align with the global interest, thanks to a "Brussels effect" where a global monopoly platform will improve privacy protection across all its markets in response to a policy change in one country. The alignment of unilateral and multilateral incentives reduces the need for a trade agreement to cover privacy protection. However, countries do have a beggar-thy-neighbor motivation to apply ad tech taxes, making these policies an area where international cooperation is needed.

JEL-Codes: F100.

Keywords: trade policy, trade agreements, WTO, platforms, two-sided markets.

Phillip McCalman
Department of Economics
University of Melbourne / Australia
mccalman@unimelb.edu.au

June 2023

1 Introduction

When the WTO launched in 1995 there were approximately 100 million computer users. While this may seem like a large number it falls well short of 1 percent of the global adult population at the time. This figure is now approaching 5 billion (approximately 80 percent of the adult population), an indication of the fundamental transformation of the global economy from one which was almost exclusively analog to one which is increasingly digital.[1] This switch has brought numerous changes, and has played a critical role in facilitating the emergence of global value chains, leading to a second unbundling (Baldwin, 2016). In this new information intensive world, the policy prescription for a country to succeed is to develop supportive infrastructure and remove barriers to the flow of information/data across borders (Ferencz, 2019). It is natural then when it comes time to negotiate trade agreements, a new core principle is the free flow of information/data between countries to aid the physical trade in goods, see Wu (2017) and Meltzer (2019).[2]

However, this physical trade facilitation motivation for policy neglects the remarkable changes that have occurred with the introduction of new virtual services. In particular, the creation of search engines and social media platforms have generated large welfare gains and proved to be highly profitable for the firms behind them, Google and Facebook. A distinctive feature of these services is their free provision, which begs the question; why are Google and Facebook so valuable? The answer lies in their two-sided business model: they take the information generated by people using their services (search queries or links/likes) and transform it into a service for third parties interested in knowing more about potential customers.

By far the most lucrative product is ad tech, which targets advertisements at specific consumers when they are likely to be at their most receptive to the idea of purchasing a

[1]See Evans (2021).

[2]Staiger (2021) examines whether e-commerce (i.e adding a digital component to the search for, order & payment of, or the delivery of goods or services) changes the nature of trade agreements when markets are competitive (i.e. no free goods or two-sided markets). In a setting where there is a trade-off between a beneficial role for digital services facilitating international trade and an associated non-pecuniary externality, he finds that the current design of the WTO is well equipped to handle e-commerce when products are traded within traditional market structures. Chen et al. (2021) find a stronger case for international coordination of data policies when privacy preferences differ across countries and a multinational uses one standard to serve all markets.

product or being persuaded on an issue. An ability to successfully provide these services has seen ad tech revenues rise from virtually zero in 2003 to two thirds of global advertising expenditure (or $US600 billion in 2021) – with 80-90% captured by Google and Facebook, which amounts to over 10% of the value global services trade.[3] This is a remarkable transformation of the advertising industry from a relatively non-traded service dominated by domestic newspapers and national television broadcasters to a highly globalized but concentrated sector – see Figure 1. Central to this shift to online advertising is the targeting of ads to individuals by using personal information. In this case though, the free flow of data is not universally seen as a good thing.

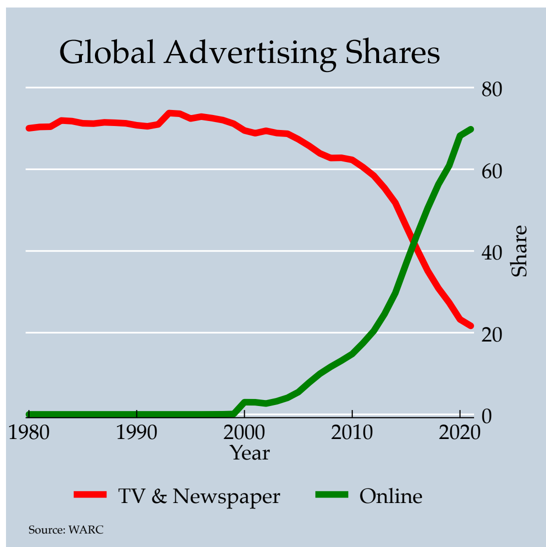

Figure 1

In particular, the implicit transaction of a free service for data/information has raised concerns over whether privacy is being protected to an appropriate extent.[4] These concerns are reflected in the aphorism: "If you are not paying for the product, you are the product," and give rise to a fear of "surveillance capitalism" more generally, see Zuboff (2019).[5] Policies put forward to counter or correct privacy issues range from the elimination of ad tech

[3]Since these revenues flow to foreign affiliates (mode 3 trade), it is not captured in the standard measures of services trade.

[4]Apart from using personal data for ad tech, there are also issues relating to the ability to protect stored details from hackers.

[5]Recent changes to WhatsApp privacy policy (owned by Facebook) saw a jump in users on Signal, a rival service that touts end-to-end encryption, from 20 million to 40 million.

through taxation to form a more conventional one-sided market (Romer, 2019), to encouraging the sharing of even more data by treating "data as labor", with tech firms paying to access it (New York Times (2019) and Arrieta-Ibarra et al. (2018)). These competing proposals reflect a more general blurring of the distinction between an economic agent as a consumer and employee, with the interactions that take place on these digital platforms having features of both roles – a theme we develop below.

While conversion of a platform to a one-sided market through taxation or treating data like labor require large structural changes in the way these markets function, actual policy initiatives have been more incremental, but have nevertheless shaped platform behavior. For example, the EU's General Data Protection Regulation (GDPR) covers data management and privacy – including consent and removal of consent rights. An illustration of the potency of these laws is the speed with which data of UK residents will be removed from Irish servers to avoid the GDPR, see The Economist (2021). The US approach has focused less on privacy and more on competition issues with the FTC launching investigations into Facebook and Google, see Financial Times (2020).[6] Smaller countries like Australia and Spain have also tried to support local journalism with laws aiming to redirect ad tech revenues from Google and Facebook to local media producers. The response from Google and Facebook has been to remove or modify service to these markets, actions which are viewed as disproportionately punitive, Financial Times (2021b).

In response to these initiatives critics argue that policy-makers don't understand the unique structure of these markets and that regulations/taxes are likely to undermine one of the most dynamic sectors of the economy.[7] This disconnect was evident at a US government hearing where Senator Orrin Hatch asked Facebook CEO Mark Zuckerberg, "How do you sustain a business model in which users don't pay for your service?" "Senator, we run ads," Mr. Zuckerberg responded. The concern is that a poor understanding of these details will inevitably lead to poorly designed policy – threatening the large R&D budgets of big tech, Varian (2021), or it stifling the vibrant tech start-up scene, or both.[8]

[6] See Maskus (2018) for an overview for the different attitudes to privacy between the EU and the US.

[7] See Sokol and Van Alstyne (2021). This concern is also echoed in official reports more generally, see Crémer et al. (2019).

[8] Amazon, Alphabet/Google were ranked number one and two in the world by firm level R& D expenditure, while Facebook ranked number 14 – see Statista, "Companies with the Highest Spending," 2018.

The portrait that emerges from this discussion is of an industry with a complex set of competing motives that raise a range of challenging questions. Do platforms have the correct incentives to protect the privacy of their users? If they don't, what is the nature and extent of the market failure? If a market failure exists, how effective are the various policy options? Will intervention to enhance privacy protection undermine innovation and the quality of service offered by a platform? Will policy differentially effect big tech and start-ups? While all of these questions can be asked at a global level, they also have analogues at the national level. Which prompts the question: are unilateral and global incentives aligned? If so, how closely. If not, how divergent are they, or does it depend on country characteristics and/or the particular policy instrument?

The goal of this paper is to provide answers to these questions. It does so by accounting for one of the key criticisms noted above: developing a model of online search and social media. The starting point is a global monopoly platform model with two potential sources of revenue: the typical one associated with the direct consumption of the monopoly product and a second derived from the number of consumers served by the platform. To fix ideas think of this second source as advertising revenue or ad tech in our context. We allow the revenue per person to be increasing in the amount of information shared/exploited by the platform to offer a service to third parties. However, greater sharing or exploitation of personal information lowers an individual's net benefit from using a platform. This tension typically results in the platform protecting privacy to some extent. Whether this trade-off is resolved in a socially desirable manner is a question we address. We also allow the platform to expend resources to improve the service it offers to consumers. This set-up is reminiscent of the Spence (1975) quality choice model, except allowing for two dimensions of choice (privacy and quality) and adding the two-sided market structure from McCalman (2020).[9]

Incorporating both the choice of privacy protection and quality in a two-sided platform model allows us to address the questions raised above. However, there is an additional feature that must be accounted for: the extreme pricing structure that sees online search

[9]The basic structure is inspired by the closed economy models of two-sided markets pioneered by Rochet and Tirole (2003), Rochet and Tirole (2006), Armstrong (2006), Hagiu (2009) and Weyl (2010).

and social media services offered for free with ad tech as the only source of revenue. Our explanation for this is both simple and novel, and is due to the differential tax treatment of positive and negative user prices. In particular, if some consumers have a negative willingness to pay (which must be the case if the market isn't covered at a zero price), then a platform may find it optimal to pay them (and the infra-marginal users) to use the platform. If an income tax is levied on negative prices (wages), then the platform perceives the marginal willingness to pay as the amount required to induce the marginal user on to the platform *plus* the associated income tax.[10] This naturally generates a kink in the willingness to pay function at a price of zero. More importantly, it generates a jump or gap in the marginal revenue function. Such a gap induces a free service outcome for a non-trivial set of the parameter space, providing an explanation for the robustness of the free service model not only across countries but also through time.

A robust implication of the free service model is that the platform will over-exploit personal information and under provide quality relative to the efficient outcome. This may seem surprising since Spence (1975) finds that a monopoly can either over or under invest in quality. The key difference is the two-sidedness of the business model and in particular the price stickiness at zero for online search and social media. This effectively transforms the platform into a price-taker on the consumption side. Consequently, the incentives faced by the platform when choosing privacy protection and quality don't reflect changes in consumer surplus, just changes in ad tech revenue. As a result, they over exploit personal information to gain additional ad tech revenue and under value the benefits of improving product quality.[11]

Having distortions on both margins puts a global policy-maker squarely in a second best world where tackling one market failure can potentially exacerbate the other. This logic lies at the heart of claims that policy aimed at improving privacy protection will undermine the incentive to invest in quality. However, we should not be too quick to accept this argument. In fact, the model predicts the opposite can be true: stronger privacy pro-

[10]This is not just a thought experiment. Microsoft offers rewards to use Bing, and when these rewards exceed $600, they will send you a 1099 form documenting this as taxable income.

[11]Alternatively, a more ad hoc zero price constraint can be imposed to gain similar results for the product market aspects of the transaction, but this approach neglects the potential "data as labor" dimension of the interaction.

tection can increase the incentive to provide quality.[12] The ability of a single instrument to reduce distortions along both margins means that it is unambiguously welfare improving. However, welfare improvements can arise even if greater privacy protection does reduce quality. We find this to be the case when a privacy policy is imposed on a Big Tech platform. In particular, given the structural features of Big Tech, there is almost no chance that higher privacy standards will lower global welfare – making privacy policy a very robust instrument when targeting these firms. In contrast, start-ups rely relatively more on quality and exploit personal information less intensively. This makes a privacy policy that targets a start-up likely to have a more pronounced impact on quality, potentially lowering welfare. Privacy policy in this case is then much harder to get right than that aimed at Big Tech – suggesting privacy policy should be applied differentially based on size.

These arguments hold when taking a global perspective. However, policy has so far been conducted or proposed at the national level. Does the division of the world into sovereign states matter? We find that if a country is large, in the sense that its policy can affect quality choice, then global and unilateral incentives are aligned for marginal changes in privacy protection, especially concerning Big Tech. An interesting feature of a large country's privacy policy is that the platform will endogenously extend it, at least partially, to all other countries. This is reminiscent of the "Brussels effect" where EU regulations diffuse globally through firm behaviour, Bradford (2012). This is often attributed to a fixed cost of product design that reduces the incentive to offer different versions across countries, Bradford (2020), Grossman et al. (2021), Parenti and Vannoorenberghe (2022) and Maggi and Mrazova (2023). However, our setting has no such tailoring costs.[13]

If unilateral incentives of large countries are typically aligned with global incentives, is there any role for a trade agreement that covers privacy policy or information flows? At least initially, the answer appears to be no. However, as noted above, there are a number of proposals to use digital taxes to achieve a range of objectives. Unlike privacy policy, international cooperation is required to align incentives with a global planner. In this case

[12] See Adjerid et al. (2016) for evidence that greater privacy protection improves innovation/quality in the health care industry.

[13] Also see Maggi and Ossa (forthcoming) for a discussion of deep trade agreements in the presence of lobby groups.

both large and small countries have a greater incentive to apply taxes than exists at the global level due to rent shifting beggar-thy-neighbor motives. This makes digital taxation an area where international cooperation is likely to be most beneficial.[14]

To establish these results section 2 presents a simple model of two-sided markets for an integrated global economy, characterizing platform behavior with respect to privacy and quality in a free service environment (section 3) and how far this diverges from efficiency (section 4). This framework is then extended in section 5 to a world fragmented into sovereign nations, where one country hosts the platform and all countries have consumers and third-parties advertisers. The incentives to choose privacy policy and the role of country size is then considered. Finally, the use of taxes is examined in section 6.

2 Model

The economic environment consists of three groups of agents; a platform, consumers and third-party advertisers. The platform provides the "infrastructure" critical for an interaction between a consumer and an advertiser. This infrastructure is assumed to be software (online search engine or social media). Given the winner-take-all dynamic in these sectors (due to network externalities or advantages from data enabled learning), we assume that the platform is a global monopoly provider of the software necessary to access the platform (i.e there might have been competition for the market but there is not competition within the market).[15]

2.1 Demand for Platform services

The demand specification is chosen to highlight issues introduced by quality choice and data sharing/exploitation. All consumers are assumed to receive a baseline level of utility when they use the platform, b. This baseline utility can be augmented by the platform

[14] See Jeon et al. (2021) and Klimenko and Qu (2023) for an analysis and discussion of policy issues that arise in an online world that focuses on language differences and e-commerce.

[15] For the role of scale in search queries see He et al. (2017). On the ability to manage a social network within the constraints of the "Dunbar number" and need for a news feed algorithm see The Telegraph (2016). For an overview of winner-take-all mechanisms for platforms see Crémer (2020), Furman et al. (2019), Stigler-Center (2019) and Crémer et al. (2019).

improving the quality of the service, q. To utilize platform services, a consumer must necessarily reveal personal information (specific search queries or links/likes on social media). The extent that the platform exploits or shares this information to offer a service to third-parties, s, tends to be viewed negatively by consumers. All three of these characteristics are taken to have a common valuation across consumers.

When it comes to their opportunity cost of using the platform, c, consumers do differ. This cost has support on $[0, \bar{c}]$ with $F(c)$ and $f(c) > 0$. Given a population of consumers, L, the number x that have an opportunity cost below c is $x = F(c)L$. It then follows that the fraction of the population with opportunity costs below c is $n \equiv F(c) = \frac{x}{L}$. For the purposes of the arguments that follow, we can simplify the analysis by adopting

$$F(c) = c^{\epsilon}, \text{ where } \epsilon > 0. \tag{1}$$

Having q and s enter the willingness to pay as linear functions implies that a platform faces marginal incentives that coincide with those of a social planner for q and s when n is fixed (more on this below). Putting this structure together, if the platform sets a price, $p \gtreqless 0$, for its services, then the marginal buyer is defined by:

$$b + q - s - c - p = 0$$

with n determined by:

$$n = \int_0^{b+q-s-p} f(c)dc = (b + q - s - p)^{\epsilon}$$

Hence,

$$p(n, s, q) = b + q - s - n^{\frac{1}{\epsilon}} \tag{2}$$

2.2 Platform

The consumers form one potential source of revenue for the platform. The second source of revenue is from the third-party advertisers. The **per person** revenue function from the cross-side mapping s and q into monetary outcomes for the platform is $r(s,q)$, and reflects the ability to target ads to an individual. Even non-targeted ads could plausibly fit this scenario – the number of people that see the display ads is the driving factor. Assume that this function is increasing in both arguments and strictly concave, $r_s > 0, r_{ss} < 0$, $r_q > 0, r_{qq} < 0$ and $r_{ss}r_{qq} - r_{sq}^2 > 0$. Notice, these conditions don't pin-down the sign of r_{sq} which can be either positive or negative – a feature which plays a role below. To further simplify the analysis we assume that the platform captures all the rents from the ad tech side of the market.[16]

Given the platform facilitates interactions through software, we take the marginal cost of an additional consumer that utilizes the platform services as zero and the marginal cost of directing an advertisement at them to be zero as well. However, the quality of consumer experience on the platform does require resources. To reflect this, we assume the cost of improving quality has the following properties: $C(q) \geq 0$ where $C_q \geq 0$ and $C_{qq} > 0$.

The platform's problem is then

$$\max_{n,q,s} \pi \quad = \quad p(n,s,q)n + r(s,q)n - C(q),$$

with associated first order conditions:

$$\pi_n \quad = \quad p + p_n n + r = 0 \tag{3}$$

$$\pi_q \quad = \quad p_q(n,s,q)n + r_q(s,q)n - C_q = 0 \tag{4}$$

$$\pi_s \quad = \quad p_s(n,s,q) + r_s(s,q) = 0 \tag{5}$$

The outcome for (3) illustrated in Figure 2 where p^P and n^P are the platform's choice of price and, hence, participation when q and s are chosen to maximize profits. A notable feature of this diagram is that the cross-side revenue from ad tech acts as a negative marginal cost of gaining another user. An additional insight is that the efficient outcome is denoted by 1 in the figure. Notice that this occurs when the willingness to pay of the marginal user is negative (thus "data as labor"), highlighting the dual nature of participation in these

[16]The ability of Facebook and Google to capture a large share of the surplus in the online advertising market is the subject of anti-trust cases in the US, EU and UK.

market as having product/factor market features.

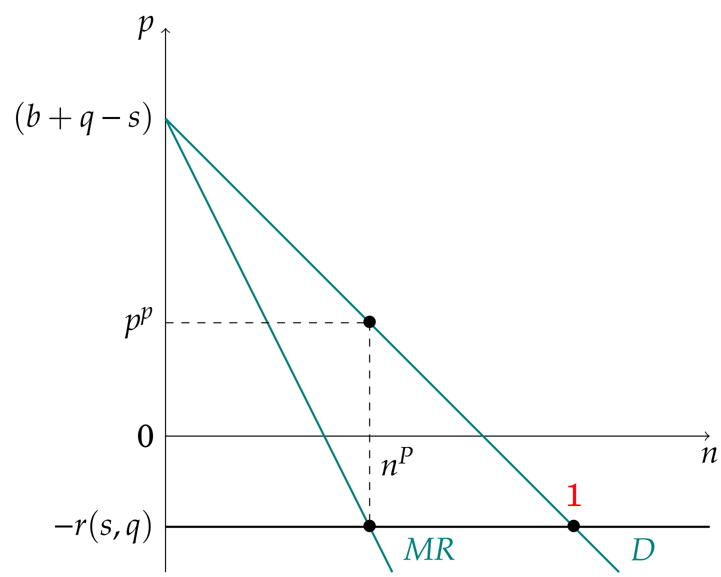

Figure 2: Two-Sided Structure

To help understand the properties of the platform's behavior, consider the efficient choices where welfare is defined as the sum of consumer surplus, advertising revenue less the cost of quality improvements:

$$W = \int_0^n p(z, s, q)dz + r(s, q)n - C(q)$$

$$\begin{aligned} W_n &= p + r = 0 \\ W_q &= \int_0^{n(s,q)} p_q(z, s, q)dz + nr_q - C_q = 0 \\ W_s &= \int_0^{n(s,q)} p_s(z, s, q)dz + nr_s = 0 \end{aligned}$$

Using the linearity of the willingness to pay function in q and s implies both p_q and p_s are constants. Hence,

$$
\begin{aligned}
W_q &= p_q n + r_q n - C_q = 0 \\
W_s &= p_s + r_s = 0
\end{aligned}
$$

This confirms the following property of the demand system:

LEMMA 1. *Given n, the choice of s and q by the platform are efficient.*

Figures 3 and 4 illustrate this outcome by comparing the incentives of a platform and a planner over the choice of s. Figure 3 shows the platform's optimal choice of s balances the loss in direct service revenue against the increase in ad tech revenue at the margin. Figure 4 shows the efficient s trades off the loss in consumer surplus against the gains in ad tech revenue. The linearity of the willingness to pay function in s ensures that the marginal revenue from direct service for s matches the marginal consumer surplus, given n. An advantage of this linear formulation is that it provides a neutral benchmark when considering the platform's incentive to choose s and q. Any other specification would result in a "Spence distortion" where the platform's incentive could lead to either over or under provision of s and q relative to the efficient level for a given n, see Spence (1975).

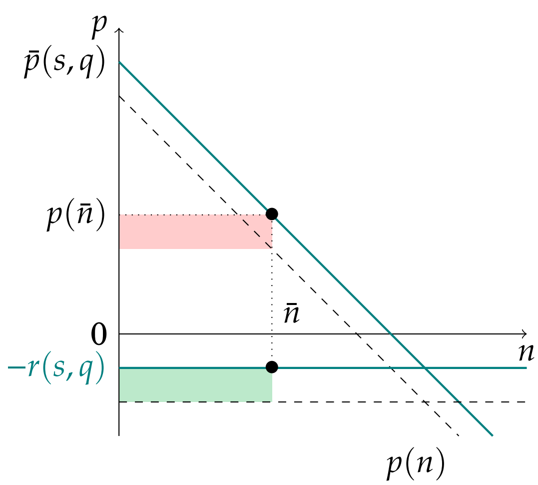

Figure 3:
Platform choice of s given $\bar{n}$

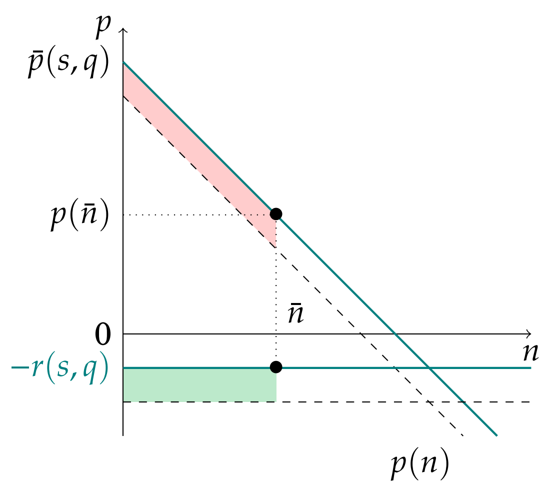

Figure 4:
Planner's choice of s given $\bar{n}$

3 Free online search and social media services

A prominent characteristic of online search and social media is that they offer their services to consumers for free. This behavior is robust through time and across countries. While such an outcome is possible in the framework developed so far, it only arises in very particular circumstances, contrary to the observed robustness. What makes a zero price so sticky? One novel possibility we explore is based on the differential treatment of positive and negative prices by a tax authority. In particular, a positive price for platform services is likely to be associated with a consumption tax, while a negative price will attract an income tax. This difference changes how the willingness to pay is perceived by the platform, generating a kink in the perceived willingness to pay function at zero. The associated marginal revenue function then has a jump or gap at the corresponding n. If the platform's profit maximizing price is zero, then this same price will be selected for a number of shocks to r, s and q.

To explore this behavior, we assume (for now) that the consumption tax rate is zero and the income tax rate, τ, is positive. In this case, to induce a "consumer" to use the platform when their opportunity cost is greater than the benefits of the platform, a payment is required to cover the difference (we'd typically think of this as a wage). In the presence of an income tax, this payment would be accompanied by τp in tax revenue. Moreover, since we focus on linear price structures, this payment would be extended to every "consumer" on the platform. Hence, the willingness to pay function (average revenue/expenditure function) has the following form:

$$\begin{aligned} p(n,s,q) \quad & \text{if } p \geq 0 \\ p(n,s,q)(1+\tau) \quad & \text{if } p < 0 \end{aligned}$$

where τ is the marginal income tax rate. This kink at $p = 0$ translates into a gap in the marginal revenue/expenditure function:

$$p + n\frac{dp}{dn} \qquad \text{if } p \geq 0 \tag{6}$$

$$\left(p + n\frac{dp}{dn}\right)(1+\tau) \qquad \text{if } p < 0 \tag{7}$$

These functions are shown in Figure 5.

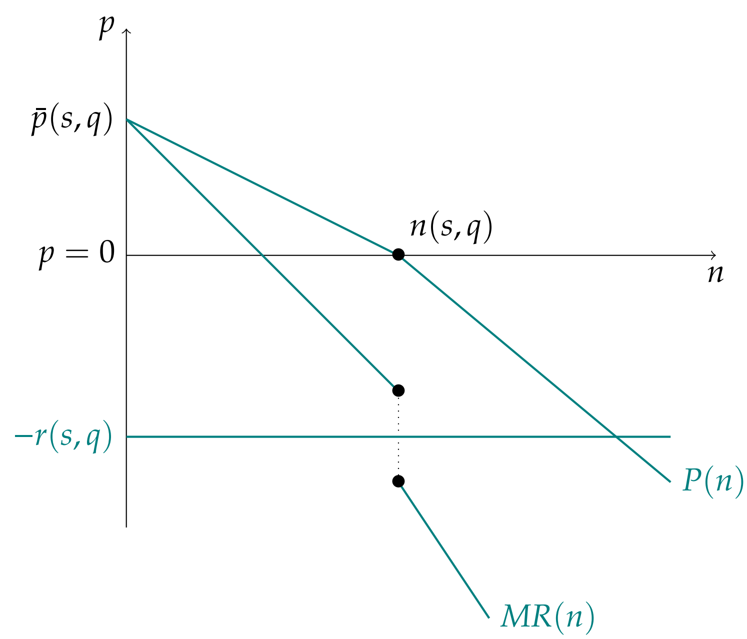

Figure 5: $p(n)$ kinked due to income taxes generates a jump in $MR(n)$.

Define the size of the gap as (6) - (7) evaluated at $p = 0$ or

$$G \equiv -n\tau \frac{dp}{dn}\bigg|_{p=0} = \frac{(b+q-s)\tau}{\epsilon}$$

This has the following properties:

LEMMA 2. *The gap in the platform's marginal revenue/expenditure function, G, is*

i. decreasing in s,

ii. increasing in q,

iii. decreasing in ϵ.

Notice in particular that higher q increases both the size of the gap and the per person ad tech revenue. This positive correlation tends to maintain the zero price outcome. In contrast, higher s produces a negative correlation between the size of the gap and r, making a deviation from the zero price outcome in the direction of a negative price more likely. We will concentrate our analysis on empirically relevant outcomes where, at least initially, the following holds:

$$\frac{(1+\tau)(b+q-s)}{\epsilon} \geq r(s,q) \geq \frac{(b+q-s)}{\epsilon} \tag{8}$$

3.1 Platform's choice of s and q with free services

Since free provision is a relatively robust outcome, it then becomes meaningful to ask: What implications does the free service model have for the platform's choice of s and q? Under this pricing configuration the platform's objective function becomes:

$$\max_{\{s,q\}} \pi \quad = \quad n(s,q)r(s,q) - C(q),$$

with the following first order conditions:

$$\pi_s \quad = \quad n(s,q)r_s(s,q) + n_s(s,q)r(s,q) = 0 \tag{9}$$

$$\pi_q \quad = \quad n(s,q)r_q(s,q) + n_q(s,q)r(s,q) - C_q = 0 \tag{10}$$

where s^P and q^P are the choices that satisfy these equations.

To evaluate social welfare at this point consider,

$$W(s,q) = \pi(s,q) + CS(s,q)$$

The marginal welfare from information sharing and quality are:

$$W_s = \pi_s(s,q) + CS_s(s,q) = \Big|_{\{s^P,\ q^P\}} \overbrace{\pi_s(s^P,q^P)}^{=0} - n(s^P,q^P) < 0 \tag{11}$$

$$W_q = \pi_q(s,q) + CS_q(s,q) = \Big|_{\{s^P,\ q^P\}} \overbrace{\pi_q(s^P,q^P)}^{=0} + n(s^P,q^P) > 0 \tag{12}$$

It then follows:

PROPOSITION 1. *When a platform maximizes profits under a free service model, personal information is over exploited (s is too high) and quality is under provided (q is too low).*

A global planner finds that private data is over shared/utilized and platform quality is too low. This holds despite the linear set-up (and Lemma 1) since the outcome occurs at a kink and effectively makes the platform a price taker – notice that there is nothing special about zero in this sense but the differing tax treatment of positive and negative prices must generate a kink at a zero price. The low quality and over-exploitation of data is also the basis for arguments that have been made against both Google and Facebook – see Morton and Dinielli (2020a) and Morton and Dinielli (2020b).[17]

4 Global Privacy Policy

A prima facie case for policy intervention follows from the presence of distortions on both the privacy and quality margins under the free service model. However, determining what tools are feasible/effective is a more complicated matter. For instance, it doesn't seem plausible to regulate quality in an industry where the pace of technological change is rapid, and even the definition of quality itself is changing. If policy toward platform quality is ruled out, can an intervention that targets privacy protection alone be welfare increasing? Given the second best nature of the setting, the answer to this question isn't obvious and there is always a concern that reducing one distortion risks exacerbating the other.

[17] More broadly the rebrand of Facebook to Meta has been interpreted as a strategy to escape the quality issues associated with Facebook and its relentless pursuit of profits over the safety of its users – see Financial Times (2021a).

To make progress, assume that a policy maker can set the degree of information sharing/exploitation by the platform at some level $\bar{s}$. We specify this policy instrument in relatively abstract terms and focus on the question of whether its possible for a policy maker to improve welfare in this single instrument setting, leaving aside the issue of exactly how this might be achieved in practice.[18]

If the policy maker selects $\bar{s}$, the platform is still free to choose quality. Hence, the platform's quality reaction function has the following form, $q^P(\bar{s})$, where the superscript P denotes that quality is chosen as a constrained optimum for the platform. Using the platform's reaction function, the planner's problem becomes:

$$\max_{\bar{s}} W(\bar{s}, q^P(\bar{s})) = \pi(\bar{s}, q^P(\bar{s})) + CS(\bar{s}, q^P(\bar{s})) \tag{13}$$

The derivative with respect to $\bar{s}$ is then:

$$W_{\bar{s}}(\bar{s}, q^P(\bar{s})) \quad = \quad \pi_{\bar{s}}(\bar{s}, q^p(\bar{s})) + CS_{\bar{s}}(\bar{s}, q^P(\bar{s}))$$

To address the question of whether privacy policy is capable of improving welfare, evaluate this derivative at the unregulated outcome, $\bar{s} = s^P$ and $q^P(s^P)$:

$$\begin{aligned} -W_{\bar{s}}(\bar{s}, q^P(\bar{s}))\Big|_{\{s^P,\ q^P\}} &= -\left(\overbrace{\pi_{\bar{s}}(s^P, q^P)}^{=0} + CS_{\bar{s}}(s^P, q^P)\right) \\ &= n(s^P, q^P)\left(1 - \frac{dq}{d\bar{s}}\right) \end{aligned} \tag{14}$$

Based on this expression we describe privacy policy as robust if $\frac{dq}{ds}$ is in the neighborhood of 0. In this case, small changes in privacy policy unambiguously improve welfare. In contrast, when $\frac{dq}{ds}$ is in the neighborhood of 1, privacy policy is fragile in the sense that increases in privacy protection can't increase welfare. The key to using privacy policy effectively is knowing how likely it is we are in a robust or fragile setting – some guidance

[18] See Stigler-Center (2019) for a discussion of the issues raised by the "how" dimension of this question.

on this is given below.

Equation (14) reflects the arguments and concerns that have been expressed more broadly about privacy protection. In particular, when starting at an unregulated outcome, the capacity of policy to improve welfare depends on whether the direct gains from less exploitation of personal information are offset by the potential change in quality. If quality is relatively responsive, then welfare can decline. This leads some commentators to worry that regulation will "break the internet". However, statements along these lines implicitly assume that $\frac{dq}{d\bar{s}}$ is unambiguously positive – i.e. there is an inherent trade-off between platform quality and the exploitation of personal data. The value of using a model is that we don't need to make such an assumption as we can directly evaluate the properties of this derivative.

Totally differentiate (10):

$$\frac{dq}{d\bar{s}} = -\frac{\pi_{q\bar{s}}}{\pi_{qq}}$$

Since $\pi_{qq} < 0$, $\frac{dq}{d\bar{s}}$ has the same sign as π_{qs}.

$$\begin{aligned} \pi_{q\bar{s}} &= n_{q\bar{s}}r + n_q r_{\bar{s}} + n_{\bar{s}} r_q + n r_{q\bar{s}} \\ &= -n_{qq}r - n_q(r_q - r_{\bar{s}}) + n r_{q\bar{s}} \end{aligned}$$

where the second line uses the symmetry (up to a sign) of s and q in determining n. Since the sign of each term is ambiguous, the overall sign is also generally ambiguous. The possibility of a negative sign offers an important result: There might not be a trade-off between privacy and platform quality. Indeed, this possibility cannot be ruled out in the framework we have adopted.

LEMMA 3. *There exists parameter configurations that satisfy $\epsilon > 0$ while ensuring that $r(s, q)$ is strictly concave and $C(q)$ is strictly convex and allow either $\frac{dq}{ds} = 0$ (robust privacy policy) or $\frac{dq}{ds} = 1$ (fragile privacy policy) to hold for an unregulated free service equilibrium.*

These results are summarized in the figures below where the two first order conditions are represented in (s, q) space. The slope of each function is $\frac{dq}{ds}\Big|_{\pi_s=0} = -\frac{\pi_{ss}}{\pi_{sq}}$ and $\frac{dq}{ds}\Big|_{\pi_q=0} =$

$-\frac{\pi_{qs}}{\pi_{qq}}$, and the relative slopes are determined by the second order condition, which can be rewritten as: $\frac{\pi_{qq}}{\pi_{qs}}\frac{\pi_{ss}}{\pi_{sq}} > 1$. Consequently, $\pi_s = 0$ always has a more pronounced slope (with both slopes having the same sign). The intersection of $\pi_s = 0$ and $\pi_q = 0$ defines the profit maximizing choices and the associated iso-welfare function is represented by $\overline{W}$, with more preferred outcomes above and/or to the left of this point. Figure 6 illustrates the limitations of a single policy instrument when $\overline{W}$ is tangent to $\pi_q = 0$ at the initial outcome (i.e. $\frac{dq}{ds} = 1$). Since any decrease in $\bar{s}$ constrains the outcome to be a movement along $\pi_q = 0$, q must be lower and so must welfare. In contrast, Figure 7 demonstrates that a trade-off between s and q need not be welfare reducing if the iso-welfare curve is steeper than $\pi_q = 0$. Finally, Figure 8 illustrates a situation with no trade-off between s and q, and improved privacy protection induces the platform to offer a higher quality.

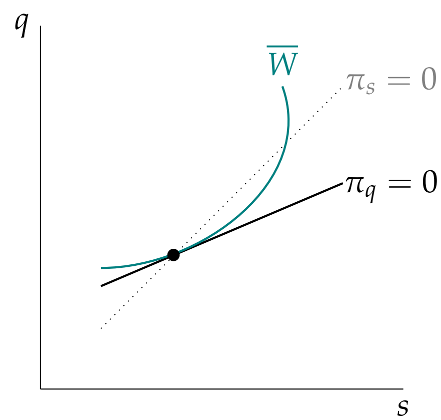

Figure 6:
No Gains and trade-off

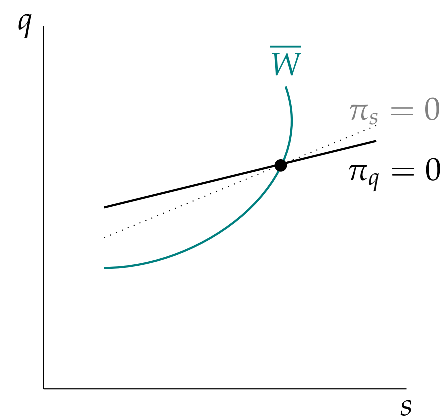

Figure 7:
Gains and trade-off

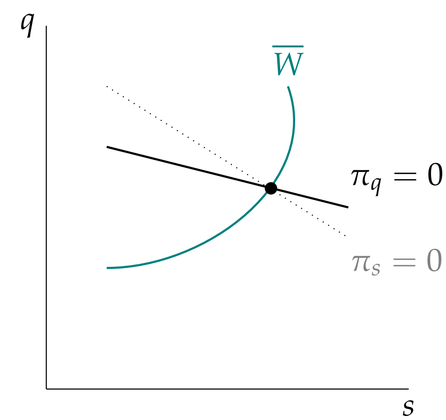

Figure 8:
Gains and no trade-off

4.1 Benchmark Model

To gain additional insight and help build intuition we construct a parametric model that embodies the possibilities outlined in Lemma 3. In particular, we define the benchmark model as having a quadratic cost of quality, a uniform distribution of the opportunity cost c and that ad tech revenue has diminishing returns in q and s which is parameterized by $\rho \in [0, 1]$:

Definition 1. *Benchmark Model:* $C(q) = kq^2$, $k > 0$, $\epsilon = 1$, $r_{qs} = 0$ *and* $r = s^\rho + q^\rho$.

The essential elements of this model take a neutral position on both the curvature of the demand function ($n_{qq} = 0$) and the cross effects within the ad tech revenue function ($r_{qs} = 0$). Under these assumptions, the demand system is linear and the ad tech revenue function is additive. Despite these restrictions, this framework is rich enough to reflect the uncertainties that exist over whether privacy policy is robust or fragile but also simple enough to allow for analytic solutions to be derived. It also enables us to examine an important feature of the tech sector, the pronounced asymmetry in firm size/age. In particular, we explore the interaction between firm size/age and the welfare effects of privacy policy.

4.1.1 Big Tech versus Start-ups

A prominent feature of the tech sector is its dichotomous firm age/size structure. For any high profile market segment there is typically one dominant firm which is characterized as "Big Tech".[19] The other type of firm is a "start-up" which typically possesses a rough/raw idea or concept and the firm itself (and its product) is not well known more generally. We distinguish between these two class of firms based on their baseline value to a consumer, which is interpreted as reflecting installed base or other accumulated data based advantages. To capture the essence of each these generic types, we set $\boldsymbol{b} = \mathbf{0}$ for a start-up and $\boldsymbol{b} > \mathbf{0}$ for Big Tech (given $n \in [0, 1]$ we have in mind $b > \frac{1}{2}$). This translates to a start-up confronted by $n = q - s$ and Big Tech facing $n = b + q - s$, when considering free service models.

Based on these definitions we address two questions for each type of firm.

1. When is privacy policy costly to welfare? (i.e. $\frac{dq}{ds} \geq 1$)
2. When is privacy policy robust (i.e. $\frac{dq}{ds} = 0$)

As discussed above, the first question provides an indication of how hard it is to get privacy policy right, while the second question says something about how hard it is to get the policy wrong. The answers to these questions prove to be especially informative. As a first step we note:

[19] Big Tech is also protected from competition by "kill zones" that reflect the idea that no funding is available from VC's for start-ups that try to compete directly with these dominant firms.

LEMMA 4. *Consider the benchmark model.*

i. When $\frac{dq}{ds} = 0$ *then* $q = s$*.*

ii. When $\frac{dq}{ds} = 1$ *then* $q > s > b$*.*

Using these properties we can now state:

PROPOSITION 2. *Consider a start-up* ($b = 0$) *in the benchmark model. When* $\rho \geq \frac{1}{2}$ *then* $\frac{dq}{ds} \geq 1$ *and privacy policy must weakly reduce welfare. Moreover, privacy policy is never robust when applied to a start-up.*

The first part of this proposition outlines a broad set of the parameter space where privacy policy will reduce welfare when applied to a start-up. The second part of the proposition reveals that any privacy policy applied to a start-up must involve a trade-off and in this sense is never robust. The intuition follows from a start-up's need to attract users through quality of service, which means it also needs to be careful not to over exploit the data it collects from users. A policy that restricts privacy beyond the start-ups choice comes at the cost of quality – privacy and quality are complements. While such a relationship seems sensible, it isn't reflected in the design of the GDPR, with a disproportionate impact on start-ups – see Janssen et al. (2021).

Given the level of uncertainty surrounding the likely success and approximate value of key parameters of a start-up, the ability to design a finely calibrated policy measure that would be required to deliver a welfare gain is also doubtful. This proposition suggests the rule of thumb that start-ups should not be specifically targeted by privacy policy. However, Big Tech is another matter.

PROPOSITION 3. *Consider Big Tech* ($b > 0$) *in the benchmark model. Privacy policy is weakly robust (i.e.* $\frac{ds}{dq} = 0$ *or* $q = s$*) when* $b \leq 2\rho^{\frac{\rho}{1-\rho}}$ *and* $k = 2^{2-\rho}(b\rho)^{\rho-1}$ *where* $q = s = \frac{b\rho}{2}$*. Moreover, as a member of Big Tech becomes sufficiently dominant,* $b \to 1$*, privacy policy is never welfare reducing.*

Unlike a start-up, a dominant Big Tech platform tends to leverage their advantages (high b due to a large market share) by more intensively exploiting personal data relative to investing in quality. Consequently, a loss at the margin from greater privacy can be partially offset by a marginal increase in platform quality – making quality and sharing substitutes. It is the relative reliance of Big Tech on the intensive use of personal data that helps identify privacy policy as welfare improving. This makes privacy policy aimed at the most dominant members of Big Tech very compelling for a policy maker.[20]

5 Platforms and Sovereign Nations

The preceding analysis takes a global perspective with a single government. In reality policy is determined by sovereign states, which we assume have the objective of using policy to improve the welfare of their citizens. The winner-take-all nature of platform competition that motivates our global monopoly platform approach also implies pronounced asymmetries across countries depending on whether or not the platform is "located" within their borders. To capture this feature we consider two regions, home and foreign. Home has the characteristic common to most countries in the world, being composed of consumers and third-party advertisers but not the platform. We maintain the assumption that third-party advertisers have their full surplus extracted by the platform. Hence, the welfare of the home region consists exclusively of consumer surplus. All variables corresponding to the home region will be indexed by h. The foreign country, on the other hand, houses all three types of actors and counts the profits of the platform as part of its welfare along with the consumer surplus of its residents (its variables will be indexed by f). Finally, consumers and third-party advertisers only interact within national boundaries, not across them, matching the observed nature of interaction.[21] To facilitate the analysis, we split the global population evenly between the two regions (asymmetries are considered below).

[20]The EU's proposed Digital Markets Act (DMA) explicitly differentiates according to firm size. In particular, the proposed rules of conduct will apply to any tech firm with a market capitalization over €80bn.

[21]According to Facebook's social connectedness index, the likelihood of an international link was less than 1% of a domestic one in 2017. Additionally, geographical location is a key variable used for targeting ads online.

5.1 Platform's choice of s_h, s_f and q with free service

Let's start by evaluating the platform's choices and confirming that they remain inefficient when the world is split into countries. In this world with two regions, the platform's free service objective function is:

$$\max_{\{s_h, s_f, q\}} \pi = \frac{1}{2} n_h(s_h, q) r_h(s_h, q) + \frac{1}{2} n_f(s_f, q) r_f(s_f, q) - C(q)$$

The first order conditions defining optimal choice are:

$$\pi_{s_h} = n_h \frac{\partial r_h}{\partial s_h} + \frac{\partial n_h}{\partial s_h} r_h = 0$$

$$\pi_{s_f} = n_f \frac{\partial r_f}{\partial s_f} + \frac{\partial n_f}{\partial s_f} r_f = 0$$

$$\pi_q = \frac{1}{2}\left(n_h \frac{\partial r_h}{\partial q} + \frac{\partial n_h}{\partial q} r_h \right) + \frac{1}{2}\left(n_f \frac{\partial r_f}{\partial q} + \frac{\partial n_f}{\partial q} r_f \right) - C_q = 0 \quad (15)$$

The global welfare function is defined as:

$$W(s_h, s_f, q) = CS_h(s_h, q) + CS_f(s_f, q) + \pi(s_h, s_f, q)$$

And when evaluated at the platform's profit maximizing choices,

$$W_{s_h}\Big|_{\{s_h^P,\, s_f^P,\, q^P\}} = -\frac{1}{2} n_h(s_h^P, q^P) < 0 \quad (16)$$

$$W_{s_f}\Big|_{\{s_h^P,\, s_f^P,\, q^P\}} = -\frac{1}{2} n_f(s_f^P, q^P) < 0 \quad (17)$$

$$W_{q}\Big|_{\{s_h^P,\, s_f^P,\, q^P\}} = \frac{1}{2} n_h(s_h^P, q^P) + \frac{1}{2} n_f(s_f^P, q^P) > 0 \quad (18)$$

These derivatives confirm that Proposition 1 extends to a setting with two regions. Since these outcomes also neglect consumer surplus in either location, each region has an incentive to use policy to improve national welfare.

5.2 National Privacy Policy

Starting at the initial equilibrium we ask: Are the incentives that shape home's privacy policy consistent or in conflict with global incentives? We begin by defining home's welfare as a function of its mandated privacy level, $\bar{s}_h$:

$$W_h(\bar{s}_h) = CS_h(\bar{s}_h, q(\bar{s}_h)) \tag{19}$$

The home country has an incentive to strengthen its privacy protection, starting from an unregulated free service equilibrium, if the following derivative is positive:

$$-\frac{\partial W_h}{\partial \bar{s}_h} = \frac{1}{2} n_h \left(1 - \frac{dq}{d\bar{s}_h}\right) \tag{20}$$

To determine the impact of changing $\bar{s}_h$ on the platform's remaining choice variables, q and s_f, totally differentiate the platform's first order conditions, π_{s_f} and π_q.

$$\begin{aligned} ds_f \pi_{s_f s_f} + dq \pi_{s_f q} &= -ds_h \pi_{s_f \bar{s}_h} (= 0) \\ ds_f \pi_{q s_f} + dq \pi_{qq} &= -ds_h \pi_{q \bar{s}_h} \end{aligned}$$

$$\underbrace{\begin{bmatrix} \pi_{s_f s_f} & \pi_{s_f q} \\ \pi_{q s_f} & \pi_{qq} \end{bmatrix}}_{A} \begin{bmatrix} ds_f \\ dq \end{bmatrix} = d\bar{s}_h \begin{bmatrix} 0 \\ -\pi_{q \bar{s}_h} \end{bmatrix} \tag{21}$$

Using Cramer's rule:

$$\frac{ds_f}{d\bar{s}_h} = \frac{\pi_{s_f q} \pi_{q \bar{s}_h}}{|A|} \tag{22}$$

$$\frac{dq}{d\bar{s}_h} = \frac{-\pi_{s_f s_f} \pi_{q \bar{s}_h}}{|A|} \tag{23}$$

Notice that beginning from the unregulated outcome, it must be that $\pi_{s_f q} = \pi_{q \bar{s}_h}$, and since $|A| > 0$, it follows $\frac{ds_f}{d\bar{s}_h} \geq 0$. Home's policy is endogenously globalized.

PROPOSITION 4. *"The Brussels Effect": When a large country unilaterally increases privacy protection (lowers $\bar{s}_h$), the platform endogenously strengthens privacy protection in other markets as well.*

Figures 9 and 10 confirm the operation of the Brussels effect is independent of the sign of π_{qs_h}. Starting with the case where $\pi_{qs_h} > 0$, an increase in home's privacy protection lowers the marginal profitability of quality. This exogenous shock causes $\pi_q = 0$ to shift down (i.e. the profit maximizing q is lower for an given s_f). Since the condition $\pi_{s_f} = 0$ is unaffected by the change in s_h, the new outcome involves a movement along this function, resulting in both a lower s_f and q. Alternatively, when $\pi_{qs_h} < 0$ both $\pi_{s_f} = 0$ and $\pi_q = 0$ are negatively sloped. A decrease in s_h now improves profitability of q at the margin, so $\pi_q = 0$ shifts upward and the new outcome is found by sliding along $\pi_{s_f} = 0$, so q is higher and s_f is lower. Hence, the Brussels effect is independent of the sign of π_{qs_h}.

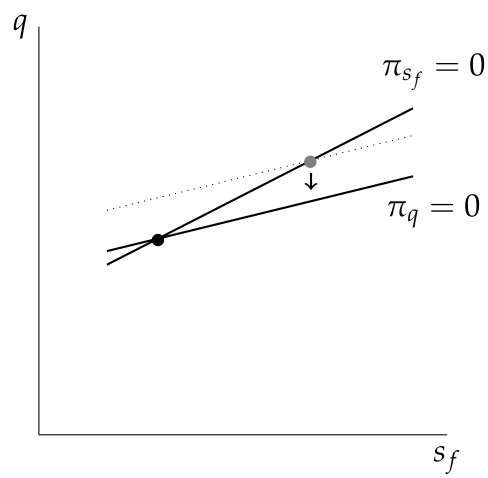

Figure 9:
Brussels effect and trade-off

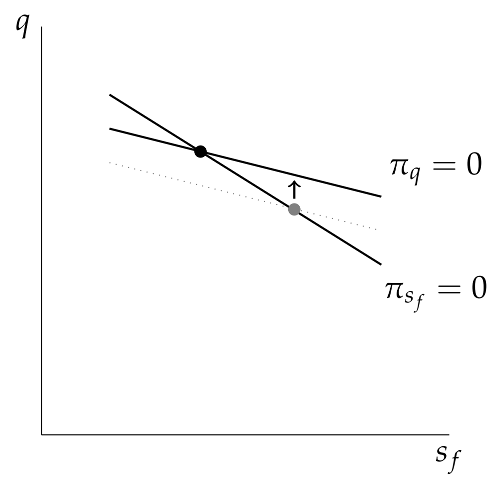

Figure 10:
Brussels effect and no trade-off

One concern when a country acts unilaterally is that its policies may not be in the global interest due to beggar-thy-neighbor incentives. Is there such a conflict for privacy policy? If so, how much tension is there between unilateral and multilateral incentives? To evaluate the alignment of home's unilateral incentives with global preferences, consider the two thresholds identified in Lemma 3. Starting with the threshold that demarcates initial outcomes where strengthening privacy policy increases quality from those that don't, $\frac{dq}{ds} = 0$, we find that the same boundary applies at the country level, $\frac{dq}{ds_h} = 0$. This result follows since we start from an initial global equilibrium where $\pi_{qs} = \pi_{sq} = 0$ (i.e. $\frac{dq}{ds_h} = 0$ implies $\pi_{qs} = \pi_{sq} = 0$). Since the platform offers the same level of privacy in both an integrated world economy and in each national economy, it must be that $\pi_{s_f q} = \pi_{s_h q} = \pi_{sq} = 0$ as well. Consequently,

LEMMA 5. *When $\frac{dq}{ds} < 0$ then $\pi_{sq} = \pi_{s_h q} < 0$ and $\frac{dq}{ds_h} < 0$, and the Brussels effect acts as a mechanism that increases welfare in the foreign country.*

The other threshold identifies initial outcomes where a global planner would not pursue increases in privacy protection. When home acts unilaterally does it pursue privacy protection beyond the global interest? This appears to be a possibility and is the outcome of two spillovers. The first is due to the Brussels effect and is always beneficial to foreign. The second spillover stems from the fact that home only considers how changes q effects its welfare. When $\pi_{qs_h} < 0$ this omission never prevents home from improving privacy protection when it's in the global interest. However, when $\pi_{qs_h} > 0$ then improved privacy protection at home lowers q, lowering foreign's welfare (negative spillover).

To evaluate the net outcome of the two spillovers consider a parameter configuration where a global planner is just indifferent about improving global privacy protection. In this case we find

LEMMA 6. *When $\frac{dq}{ds} = 1$ then $\frac{ds_f}{ds_h} < \frac{dq}{ds_h} < 1$ and the unilateral incentive to initiate an increase in privacy protection exceeds the global incentive.*

Together these two lemmas demonstrate that whether or not unilateral and global incentives are aligned depends on underlying parameter values – especially in the case of Lemma 6. We can be more explicit about the relevant parameter ranges by returning to the benchmark model. Now we can make a stronger statement:

PROPOSITION 5. *Under the benchmark model, global and national incentives to pursue greater privacy protection are always aligned in the presence of dominant a Big Tech platform (i.e. $b \to 1$ and $\frac{dq}{ds_h} < \frac{dq}{ds} < 1$).*

This proposition provides reassurance that the incentives of large countries to enhance privacy protection are not in conflict with global preferences provided privacy policy targets Big Tech. This is true despite the presence of international externalities. Moreover, both home and foreign face the same incentives, so, at least for marginal improvements in privacy protection, there is consensus over whether or not to pursue a privacy policy initiative.

6 Ad tech Taxation

The taxation of a platform's ad tech revenue has also received considerable attention, both as a source of government revenue and as a corrective measure (to curb surveillance capitalism). Starting with an integrated global perspective (i.e. no individual countries), let t be the percentage tax rate applied to ad tech revenue. The introduction of an ad tech tax modifies the platform's objective function to be:

$$\max_{\{s,q\}} \pi = \frac{n(s,q)r(s,q)}{1+t} - C(q)$$

With associated first order conditions,

$$\pi_s = \frac{n(s,q)r_s(s,q) + n_s(s,q)r(s,q)}{1+t} = 0 \tag{24}$$

$$\pi_q = \frac{n(s,q)r_q(s,q) + n_q(s,q)r(s,q)}{1+t} - C_q = 0 \tag{25}$$

The sign of $|A|$ is still determined by the same considerations as (21) and using Cramer's rule:

$$\frac{ds}{dt} = \frac{\pi_{sq}\pi_{qt}}{|A|} \tag{26}$$

$$\frac{dq}{dt} = \frac{-\pi_{ss}\pi_{qt}}{|A|} \tag{27}$$

Since $\pi_{qt} < 0$ and $|A| > 0$ we have:

LEMMA 7. *Starting at a laissez faire free service equilibrium, the introduction of a small tax on ad tech revenue unambiguously lowers platform quality, and lowers information sharing/exploitation if* $\pi_{sq} > 0$ *and increases information sharing/exploitation if* $\pi_{sq} < 0$.

Figure 11 and 12 depict these comparative statics and make clear that an ad tech tax reduces the marginal profitability of quality, $\pi_{qt} < 0$ but has no impact on the marginal profitability of data sharing, $\pi_{st} = 0$. That is, the ad tech tax shifts $\pi_q = 0$ downward, while leaving $\pi_s = 0$ in place. Since the new equilibrium is determined as a movement along $\pi_s = 0$, the slope of this curve dictates the implications for s. When $\pi_{qs} > 0$, as in Figure 11, the ad tech tax reduces both quality and data sharing. In contrast, when $\pi_{qs} < 0$, quality is reduced but the extent of data sharing is increased by the ad tech tax.

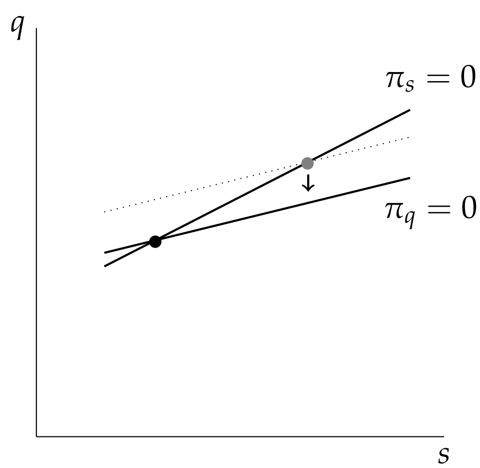

Figure 11:
ad tech tax when $\pi_{qs} > 0$

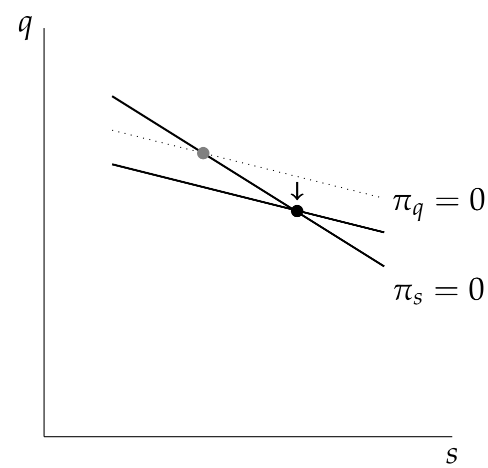

Figure 12:
ad tech tax when $\pi_{qs} < 0$

A notable implication of this lemma is that taxing ad tech doesn't guarantee that less sharing of data will occur. In fact, an ad tech tax results in more sharing in exactly the circumstances where direct regulation is most effective, i.e. when $\pi_{sq} < 0$.

To explore the welfare consequences of an ad tech tax, consider

$$\begin{aligned} W(s,q) &= \frac{n(s(t),q(t))r(s(t),q(t))}{1+t} - C(q(t)) + CS(s(t),q(t)) + \frac{t}{1+t}\Big(n(s(t),q(t))r(s(t),q(t))\Big) \\ &= n(s(t),q(t))r(s(t),q(t)) - C(q(t)) + CS(s(t),q(t)) \end{aligned}$$

Starting from an untaxed/unregulated outcome it follows that:

$$W_t\Big|_{t=0} = n\left(\frac{dq}{dt} - \frac{ds}{dt}\right) \gtrless 0 \tag{28}$$

Although the welfare consequences of an ad tech tax are generally ambiguous, we nevertheless are able to observe:

PROPOSITION 6. *If* $\pi_{qs} < 0$ *then the introduction of an ad tech tax lowers global welfare, while strengthening privacy protection raises it.*

Moreover, if we adopt the structure of the benchmark model, it immediately follows that only a Big Tech platform operates in the parameter range where $\pi_{qs} < 0$. So a global planner will choose privacy policy rather than taxation when addressing the market failures associated with Big Tech.

6.1 Large Country

When a large country evaluates the merits of an ad tech tax policy, it does so with a similar perspective to a global planner except now revenue considerations play a role.

$$W^h(s(t_h), q(t_h)) \quad = \quad CS^h(s(t_h), q(t)) + \frac{t_h}{1+t_h}\Big(n_h(s(t_h), q(t_h))r_h(s(t_h), q(t_h))\Big)$$

Starting from an untaxed/unregulated outcome it follows that:

$$W_t^h\Big|_{t_h=0} = n_h\left(\frac{dq}{dt_h} - \frac{ds}{dt_h} + r_h\right) \tag{29}$$

Notice that (24) still holds for both home and foreign, implying that s is common to both markets and (26) and (27) describe the comparative statics (with the appropriate modification that the tax is only applied in the home market). (29) reveals a "beggar-thy-neighbor" motive when a large home country taxes ad tech due to the presence of $r_h > 0$. This motive causes the unilateral and multilateral incentives to diverge over a wide range of parameter values. For example, if a global planner is indifferent about using ad tech taxes, a large home country isn't since it gains tax revenue from the foreign platform. Home's "beggar-thy-neighbor" incentive means that it will always impose a higher tax on ad tech than is efficient, providing the basis for a trade agreement to cover digital taxation.

6.2 Small Country

While "beggar-thy-neighbor" incentives are usually associated with large countries, this motive now also applies to a small country but with an added dimension. For our purposes a small country is one which takes the quality of service as exogenous. In particular, a small country has the defining characteristic that $\frac{dq}{dt_h} = 0$. In this case a small country's welfare is:

$$W^i(s(t_i)) = CS^i(s(t_i)) + \frac{t_i}{1+t_i}\Big(n_i(s(t_i))r_i(s(t_i))\Big)$$

Starting from an untaxed/unregulated outcome it follows that:

$$W^i_{t_i}\Big|_{t_i=0} = n_i r_i > 0 \tag{30}$$

where the fact that $\frac{dq}{dt_i} = 0$ implies $\frac{ds}{dt_i} = 0$ (from (24) and (25)). Consequently,

PROPOSITION 7. *Starting at a free service equilibrium, the introduction of a small tax on ad tech revenue unambiguously raises the welfare of a small country.*

More generally,

$$W^i_{t_i} = \frac{n_i r_i}{(1+t_i)^2} > 0$$

This implies a small country has a strong incentive to tax ad tech revenue. This policy could well cause a platform to re-evaluate its business model in that country, and move to a subscription service along the lines suggested by Romer (2019). This is indeed the case, as there must exist a finite ad tech tax where the platform is indifferent between the "free service" model and one where the platform offers a subscription service (i.e. $p > 0$) – see Figure 13. In particular, selecting $1 + t_i = \frac{r_i}{-p_{n_i} n_i}$ places the platform in this position. Notice that the right hand side of this equation is a constant from a small country's perspective for any outcome associated with the gap in the marginal revenue function.

This last observation raises the question of whether a small country can manipulate the location of the gap in the marginal revenue function to its advantage. In fact, it can achieve a better outcome by taxing subscription revenue with an ad valorem consumption tax of ϕ. Initially, this may seem like a strange policy to adopt given that $p = 0$ and consequently no revenue is raised. However, the objective of the small country is not to raise tax revenue but to manipulate (perceived) marginal revenue.

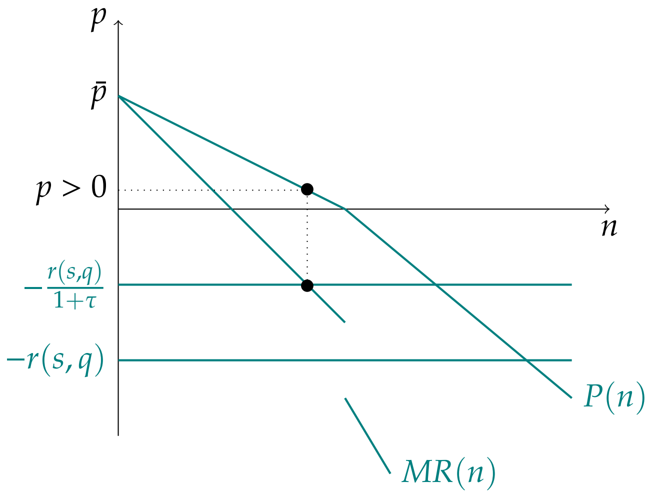

Figure 13:
ad tech tax induces $p > 0$

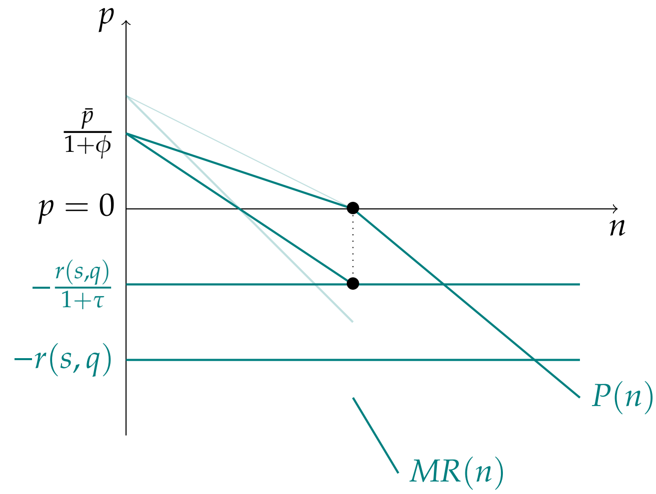

Figure 14:
Consumption tax restores initial n

To see how this works, consider the willingness to pay or average revenue function facing the platform in the presence of a consumption tax: $\frac{p(n_i)}{1+\phi}$. It follows from this that the perceived marginal revenue is $\frac{p}{1+\phi} + \frac{p_{n_i} n_i}{1+\phi}$. If we focus on the outcome where $p = 0$, this coincides with the upper bound of the gap in perceived marginal revenue, $\frac{p_{n_i} n_i}{1+\phi}$. Consequently, by increasing the consumption tax on platform services, a small country can decrease the (absolute value) of the upper bound on the perceived marginal revenue gap – see Figure 14.[22] It immediately follows that:

PROPOSITION 8. *A small country can maintain the s and n_i from the initial free service outcome and fully appropriate the surplus for its market by choosing t_i and ϕ such that $(1 + t_i) = \frac{r_i}{-p_{n_i} n_i}(1 + \phi)$ and setting ϕ arbitrarily large.*

[22]The lower bound is left unaltered by the consumption tax.

The ability to fully extract the surplus through taxation represents an extreme "beggar-thy-neighbor" mechanism.[23] Ultimately both small and large countries have an incentive to use ad tech taxes (and possibly other taxes) in a manner that departs significantly from the aims of a global social planner. This suggests any extension of trade agreements to include e-commerce should focus primarily on digital taxation and leave privacy to be determined on a unilateral basis with a focus on Big Tech.

7 Conclusion

The transformation of the economy from analog to digital technology has not only altered production structures but also created entirely new consumption experiences/services. A distinctive feature of the business models that deliver these new services is their two-sided nature. While these business models are not new, their prominence and global reach is. Gaining an understanding of what this means for welfare has been confounded by the free provision of online search and social media services. In particular, this has blurred the distinction between an economic agent as a consumer and worker/input/resource, with the interactions that take place on these platforms having features of both roles. To reflect this duality we develop a model that not only captures the two-sidedness of these markets but also accounts for the free service offered to one side. This last feature has distinctive implications for welfare, with the online search and social media platforms departing from efficiency for both privacy protection and quality of service.

These distortions provide a prima facie case for government intervention but determining how to formulate an effective policy is not obvious. Moreover, the difficulty of arriving at globally appropriate policies is increased when choices are made on a unilateral basis. The use of ad tech taxes is most prone to misuse when set by individual countries since they have a beggar-thy-neighbor dimension due to international rent shifting. Both large and small countries face these incentives. In contrast, unilateral efforts to protect privacy are less susceptible to beggar-thy-neighbor motives. While they do give rise to an international

[23]For an analysis of trade policy and market power along more traditional market structure lines and how it relates to trade agreements see McCalman (2021).

externality in the form of the Brussels effect, this typically serves global interests rather than undermining them. Together this suggests that global coordination or discipline is most needed with respect to digital taxation and that privacy policy can be left relatively unconstrained, at least for now.

A Appendix

A.1 proof of Lemma 3

To see that it is possible to construct an equilibrium where $\frac{dq}{d\bar{s}} = -\frac{\pi_{q\bar{s}}}{\pi_{qq}} = 1$, note that this implies $\pi_{qs} = -\pi_{qq} > 0$ or

$$\begin{aligned} 0 &= \pi_{qs} + \pi_{qq} \\ \Rightarrow C_{qq} &= n_q(r_s + r_q) + n(r_{qs} + r_{qq}) \end{aligned} \tag{31}$$

Since we are starting at an unregulated outcome, we also require

$$\begin{aligned} & \pi_{qq}\pi_{ss} - \pi_{qs}^2 > 0 \\ \Rightarrow \quad & \pi_{qq}(\pi_{ss} - \pi_{qq}) > 0 \\ \Rightarrow \quad & \pi_{ss} - \pi_{qq} = -2n_q(r_s + r_q) + n(r_{ss} - r_{qq}) + C_{qq} < 0 \\ \Rightarrow \quad & -n_q(r_s + r_q) + n(r_{ss} + r_{sq}) < 0 \end{aligned} \tag{32}$$

It is then straightforward to construct examples where this holds. For example, suppose $C(q) = aq^\theta$ where $\theta > 1$ and $a > 0$, while $r(s, q) = (s^\rho + q^\rho)^\alpha$ where $\rho < 1$ and $\rho\alpha < 1$.

Now let's examine $\frac{dq}{ds} = 0$ which implies $\pi_{qs} = 0$ or

$$\pi_{qs} = -n_{qq}r - n_q(r_q - r_s) + nr_{qs} = 0$$

Notice that if $n_{qq} = r_{qs} = 0$ then a $C(q)$ that delivers $r_s = r_q$ in the initial equilibrium satisfies this condition. Typically, this requires C_q to be sufficiently large.

Note that $r_{qs} = 0$ can be consistent with both outcomes and illustrates the difficulty in demarcating the parameter space in a manner that allows specific statements to be made about what determines the sign and size of $\frac{dq}{ds}$.

A.2 proof of Lemma 4

The platform's objective function in the benchmark model is

$$\pi = (b + q - s)(s^\rho + q^\rho) - kq^2$$

The first order conditions are:

$$\pi_s = -(s^\rho + q^\rho) + (b + q - s)\rho s^{\rho-1} = 0 \tag{33}$$

$$\pi_q = (s^\rho + q^\rho) + (b + q - s)\rho q^{\rho-1} - 2kq = 0 \tag{34}$$

When $\frac{dq}{ds} = 0$ we know that $\pi_{qs} = 0$ or using the above derivatives $\pi_{qs} = \rho(s^\rho - q^\rho) = 0$ which is only satisfied when $s = q$.

When $\frac{dq}{ds} = 1$ implies $\pi_{qs} = -\pi_{qq}$. Using the above derivatives:

$$\begin{aligned} \pi_{qs} &= r_s - r_q & (35) \\ -\pi_{qq} &= 2k - 2r_q - nr_{qq} & (36) \end{aligned}$$

$$\begin{aligned} \Rightarrow \quad & 2k - 2r_q - nr_{qq} = r_s - r_q \\ & 2k - (r_s - r_q) - nr_{qq} = 0 \\ & 2k - \frac{2kq}{n} - nr_{qq} = 0 \quad (\text{using } \ \pi_s + \pi_q = 0) \\ & 2\left(1 - \frac{q}{n}\right) - nr_{qq} = 0 \end{aligned}$$

Since $r_{qq} < 0$ it follows that $q > n = b + q - s > 0$, so $q > s > b$.

A.3 proof of Proposition 2

Start by rewriting the first order conditions as:

$$\begin{aligned} s^{1-\rho}\pi_s &= -s - q^\rho s^{1-\rho} + \rho(b + q - s) = 0 \\ q^{1-\rho}\pi_q &= q + q^{1-\rho}s^\rho + \rho(b + q - s) - 2kq^{2-\rho} = 0 \end{aligned}$$

$$s^{1-\rho}\pi_s + q^{1-\rho}\pi_q = (q^{1-\rho}s^\rho - q^\rho s^{1-\rho}) + (q - s) + 2\rho(b + q - s) - 2kq^{2-\rho} = 0$$

$$\Rightarrow (b + q - s) = \frac{2kq^{2-\rho} - (q^{1-\rho}s^\rho - q^\rho s^{1-\rho}) + b}{1 + 2\rho} \tag{37}$$

When $\frac{dq}{ds} = 1$ it follows that $\pi_{qs} = -\pi_{qq}$ or

$$\begin{aligned} & \rho(s^{\rho-1} - q^{\rho-1}) = -2\rho q^{\rho-1} - (\rho - 1)\rho(b + q - s)q^{2-\rho} - 2k \\ \Rightarrow \quad & \rho(s^{\rho-1} + q^{\rho-1}) + (\rho - 1)\rho(b + q - s)q^{2-\rho} - 2k = 0 & (38) \end{aligned}$$

Using the definition of a start-up, $b = 0$, and setting $\rho = 1/2$, (37) and (38) become:

$$\begin{aligned} & (q - s) = kq^{3/2} \\ & \frac{1}{2}(q^{-\frac{1}{2}} + s^{-\frac{1}{2}}) - \frac{1}{4}(q - s)q^{\frac{3}{2}} - 2k = 0 \end{aligned}$$

Substituting the first equation into the second and simplifying gives,

$$q^{-\frac{1}{2}} + s^{-\frac{1}{2}} = \frac{9}{2}k \tag{39}$$

which can be re-expressed as:

$$q^{\frac{1}{2}} = \frac{2s6\frac{1}{2}}{9ks^{\frac{1}{2}} - 1} \tag{40}$$

$$q = \frac{4s}{(9ks^{\frac{1}{2}} - 2)^2} \tag{41}$$

Evaluating $\pi_s = 0$ gives

$$s = \frac{64}{27^2 k^2}, \quad q = \frac{64}{9^2 k^2} \tag{42}$$

So $\frac{dq}{ds} = 1$ for $k \geq \left(\frac{8}{9}\right)^2$ where this restriction ensures that $n \leq 1$ in equilibrium. Moreover, if the values from (42) are substituted into (38) when $\rho \neq \frac{1}{2}$, it is straightforward to show that (38) is positive when $\rho > \frac{1}{2}$ (i.e. $\frac{dq}{ds} > 1$) while (38) is negative when $\rho < \frac{1}{2}$ (i.e. $\frac{dq}{ds} < 1$).

Finally, since Lemma 4 implies that $q = s$ when $\frac{dq}{ds} = 0$, it follows that this can only be profit maximizing for a start-up when $q = s = 0$, that is a start-up chooses not be in business.

A.4 proof of Proposition 3

Let's start by identifying the parameters that satisfy $\frac{dq}{ds} = 0$ or $q = s$ from Lemma 4. Using the first order conditions and $n = b + q - s = b$:

$$\pi_s = -2s^\rho + b\rho s^\rho = 0$$

$$\Rightarrow \quad s = q = \frac{b\rho}{2}$$

Then using $\pi_q = 0$ implies $k = 2^{2-\rho}(b\rho)^{\rho-1}$ for $b \leq 2\rho^{\frac{\rho}{1-\rho}}$ where this restriction ensures that the free-service outcome is profit-maximizing since it guarantees $r > MR = n = b \leq 1$.

Finally, for privacy policy to be welfare reducing it must be the case that $\frac{dq}{ds} \geq 1$ which requires $q > s$ from Lemma 4. However, as $b \to 1$ the requirement that $n \leq 1$ is violated and the free service outcome is no longer part of an equilibrium.

A.5 proof of Lemma 6

To evaluate this issue, consider an initial equilibrium where $\frac{dq}{ds_h} = -\frac{\pi_{s_f s_f}\pi_{qs_f}}{\pi_{s_f s_f}\pi_{qq} - (\pi_{qs_f})^2}$. Once again we can exploit the isomorphic nature of the equilibrium in an integrated world and one divided into two large countries, which implies π_{qq} is the same in both cases but $\pi_{qs_f} = \frac{1}{2}\pi_{qs}$ and $\pi_{s_f s_f} = \frac{1}{2}\pi_{ss}$. Starting from $\frac{dq}{ds} = 1$, which implies, $\pi_{qq} = -\pi_{qs}$, and noting that we also require $\pi_{ss}\pi_{qq} - (\pi_{qs})^2 = \pi_{ss}\pi_{qq} - (\pi_{qq})^2 = \pi_{qq}(\pi_{ss} - \pi_{qq}) > 0$ or $-\pi_{ss} > -\pi_{qq}$, it follows that $\frac{dq}{ds_h} = \frac{\pi_{ss}}{2\pi_{ss}-\pi_{qq}} < 1$.

A similar argument implies $\frac{ds_f}{ds_h} = \frac{\pi_{qq}}{2\pi_{ss}-\pi_{qq}}$. Since $\frac{ds_f}{dq} = \frac{ds_f/ds_h}{dq/ds_h} = \frac{\pi_{qq}}{\pi_{ss}} = \frac{\pi_{qs}}{\pi_{ss}} < 1$. Consequently, $ds_h > dq > ds_f > 0$.

References

Adjerid, Idris, Alessandro Acquisti, Rahul Telang, Rema Padman, and Julia Adler-Milstein (2016) ‘The impact of privacy regulation and technology incentives: The case of health information exchanges.’ *Management Science* 62(4), 1042–1063

Armstrong, Mark (2006) ‘Competition in two-sided markets.’ *The RAND Journal of Economics* 37(3), 668–691

Arrieta-Ibarra, Imanol, Leonard Goff, Diego Jiménez-Hernández, Jaron Lanier, and E Glen Weyl (2018) ‘Should we treat data as labor? moving beyond "free".’ *AEA Papers and Proceedings* 108, 38–42

Baldwin, Richard (2016) *The great convergence* (Harvard University Press)

Bradford, Anu (2012) ‘The brussels effect.’ *Northwestern University Law Review* 107(1), 1–69

_ (2020) *The Brussels Effect, How the European Union Rules the World* (Oxford University Press)

Chen, Yongmin, Xinyu Hua, and Keith E Maskus (2021) ‘International protection of consumer data.’ *Journal of International Economics*

Crémer, Jacques (2020) ‘The economics of the digital economy.’ WTO Trade Dialogues Lecture/Webinar Series

Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer (2019) ‘Competition policy for the digital era.’ *Report for the European Commission*

Evans, Benedict (2021) ‘The great unbundling.’ Protocol Tech in 2021, www.ben-evans.com

Ferencz, Janos (2019) *The OECD digital services trade restrictiveness index* (OECD)

Financial Times (2020) ‘Big tech faces its standard oil moment.’ December, 12

_ (2021a) ‘Facebook changes name to meta in corporate rebranding.’ October, 29

_ (2021b) ‘Facebook: content ferment.’ February, 19

Furman, Jason, Diane Coyle, Amelia Fletcher, Derek McAuley, and Philip Marsden (2019) ‘Unlocking digital competition: Report of the digital competition expert panel.’ *UK government publication, HM Treasury*

Grossman, Gene M, Phillip McCalman, and Robert W Staiger (2021) ‘The "new" economics of trade agreements: From trade liberalization to regulatory convergence?’ *Econometrica* 89(1), 215–249

Hagiu, Andrei (2009) ‘Two-sided platforms: Product variety and pricing structures.’ *Journal of Economics & Management Strategy* 18(4), 1011–1043

He, Di, Aadharsh Kannan, Tie-Yan Liu, R Preston McAfee, Tao Qin, and Justin M Rao (2017) 'Scale effects in web search.' In 'International Conference on Web and Internet Economics' Springer pp. 294–310

Janssen, Rebecca, Reinhold Kesler, Michael Kummer, and Joel Waldfogel (2021) 'Gdpr and the lost generation of innovative apps'

Jeon, Doh-Shin, Bruno Jullien, and Mikhail Klimenko (2021) 'Language, internet and platform competition.' *Journal of International Economics* 131, 103439

Klimenko, Mikhail, and Jingwen Qu (2023) 'Global digital platforms, technology transfer and foreign direct investment policies in two-sided markets.' *Economic Inquiry*

Maggi, Giovanni, and Monika Mrazova (2023) 'Harmonization...what else? the role for international regulatory agreements'

Maggi, Giovanni, and Ralph Ossa (forthcoming) 'The political economy of international regulatory cooperation.' *American Economic Review*

Maskus, Keith (2018) 'Fostering innovation in digital trade.' In *Intellectual Property and Digital Trade in the Age of Artificial Intelligence and Big Data,* ed. Xavier Seuba, Christophe Geiger, and Julien Penin (Geneva: CEIPI-ICTSD,) pp. 19–28

McCalman, Phillip (2020) 'Trade policy with fangs (aka trade policy and multi-sided platforms).' University of Melbourne

_ (2021) 'Robust trade policy to offset foreign market power.' University of Melbourne

Meltzer, Joshua P (2019) 'Governing digital trade.' *World Trade Review* 18(S1), S23–S48

Morton, Fiona M Scott, and David C Dinielli (2020a) 'Roadmap for a digital advertising monopolization case against google.' *Omidyar Network, June*

_ (2020b) 'Roadmap for an antitrust case against facebook.' *Omidyar Network, June*

New York Times (2019) 'We talked to andrew yang. here's how he'd fix the internet.' October, 15

Parenti, Mathieu, and Gonzague Vannoorenberghe (2022) 'A simple theory of deep trade integration'

Rochet, Jean-Charles, and Jean Tirole (2003) 'Platform competition in two-sided markets.' *Journal of the European Economic Association* 1(4), 990–1029

_ (2006) 'Two-sided markets: a progress report.' *The RAND journal of economics* 37(3), 645–667

Romer, Paul (2019) 'A tax that could fix big tech.' *The New York Times,* May, 6

Sokol, D Daniel, and Marshall Van Alstyne (2021) 'The rising risk of platform regulation.' *MIT Sloan Management Review* 62(2), 6A–10A

Spence, A Michael (1975) 'Monopoly, quality, and regulation.' *The Bell Journal of Economics* pp. 417–429

Staiger, Robert W (2021) *A World Trading System for the Twenty-First Century*

Stigler-Center (2019) 'Stigler committee on digital platforms: Final report.' *Stigler Center web: https://www. chicagobooth. edu/-/media/research/stigler/pdfs/digital-platforms—committee-report—stigler-center. pdf*

The Economist (2021) 'Chlorinated data: Why google and facebook are shifting british data to america.' January, 7

The Telegraph (2016) 'Facebook users have 155 friends.' January, 20

Varian, Hal R (2021) 'Seven deadly sins of tech?' *Information Economics and Policy* 54,

Weyl, E Glen (2010) 'A price theory of multi-sided platforms.' *American Economic Review* 100(4), 1642–72

Wu, Mark (2017) 'Digital trade-related provisions in regional trade agreements: Existing models and lessons for the multilateral trade system.' *ICTSD*

Zuboff, Shoshana (2019) *The age of surveillance capitalism: the fight for a human future at the new frontier of power* (Public Affairs)