The Competitive Effects of Search Engine Defaults

Erik Hovenkamp*

October 29, 2024

Abstract

The landmark antitrust case *U.S. v. Google* centers on vertical contracts in which Google pays device makers and web browsers to make Google Search the default search engine on their products. But the competitive effects of such arrangements are not yet well understood. To study them, I introduce a novel model of competition between two-sided search platforms, which earn all revenues on the advertiser side. Search algorithms "learn" and improve with use, effectively creating network effects on the consumer side. Defaults create switching costs that "nudge" consumers toward the default platform. Due to algorithmic learning, defaults can have significant competitive effects even if switching costs are small. Broad defaults (those affecting a large share of users) by a dominant platform reduce consumer welfare under most plausible conditions; they also suppress entry and investment by laggards. However, narrow defaults (particularly by laggards) can create positive spillovers and encourage entry and investment.

*Cornell University. For helpful comments, I thank Javier Donna, Josh Fischman, Rich Gilbert, Andrei Hagiu, Tom Hubbard, Justin Johnson, Mike Katz, Doug Melamed, Eric Posner, Steve Salop, Megan Stevenson, and Mike Whinston. I am also grateful to organizers and participants at the Northwestern Conference on Antitrust Economics and Competition Policy, and the UVA Law and Economics Workshop.

1 Introduction

In the landmark case *United States v. Google*, Google is accused of monopolizing the market for internet search.[1] The case centers on agreements in which Google pays other firms, such as mobile device makers (e.g. Apple and Samsung) and web browsers(e.g. Mozilla, Opera), to ensure that Google Search is the default search engine on their products. These deals cover most mobile devices sold in the United States. And about 70-80% of all search queries are made on browsers or search apps on which Google Search is the default search engine.

A user can switch the default search engine to a competing alternative (e.g. Bing) through the settings menu. For a moderately tech-savvy user, this is quite easy to do. Thus, it is tempting to conclude that default status has no meaningful impact on consumer choice, since it is easily bypassed. If this were true, default status would confer little value to a search engine. But Google pays tens of billions of dollars annually to secure default status for Google Search. Indeed, in 2021 it paid about $20b to Apple alone. For this immense price tag to make sense, defaults must be somewhat "sticky." That is, they must affect the search engine choice of an appreciable number of consumers.

In this paper, I study the competitive effects of search engine defaults. Three important features of the search market make this research problem especially rich. First, while we can usually focus on price effects to gauge effects on competition and consumer welfare, this is not possible in the search market because search engines do not charge prices to consumers. This relates to the second important feature, which is that search engines are two-sided platforms that earn all revenues on the advertising side. Finally, a search algorithm "learns" and improves as consumers use it. This effectively creates direct (within-group) network effects on the consumer side of the market, which influences competition and welfare in important ways.

To study these deals, I introduce a model of competition between two-sided search platforms, one of which is dominant. The platforms are vertically and horizontally differentiated. Due to algorithmic learning, there are within-group network effects on the consumer side of the market. The platforms charge prices to advertisers, but not consumers. They compete for consumers based on quality. Specifically, aside from its advertising price, each platform chooses a variable (which I call the "ad intensity level") that governs the extent to which the platform's advertising activities detract from the consumer experience. This generalizes a number of ideas in the literature about how search engines exercise market power over consumers without charging them prices, such as by increasing the number or intrusiveness of unwanted ads, by collecting more data about user activity, or by distorting search results in ways that benefit advertisers.

When a platform obtains default status, consumers must incur a switching cost to change to a different search platform. A default agreement may also contain restrictions that act to magnify switching costs.[2] Such deals have a biasing effect on consumer choice. It is analogous

[1]The trial court reached a decision a few months after this paper first appeared online. The court found Google liable, but Google will surely appeal that decision.

[2]For example, the Justice Department alleges that many of Google's deals with device makers prohibit "the preinstallation of any rival general search services" on their devices. This is analogous to contractual restraints in the *Microsoft* case that discouraged computer makers from preinstalling rival web browsers.

to imposing a tax on all but one of the products in a market. This "nudges" consumers toward the default platform.

So long as algorithmic learning benefits are not too strong, a default agreement involving the dominant platform will always reduce static consumer welfare. Like a tax, switching costs create frictions discourage adoption of rival search platforms. As a result, some consumers who would have opted for a rival search engine will instead stick with the default platform. Due to algorithmic learning, this exodus of users reduces the quality of rivals' search algorithms, leading even more consumers to abandon them.

A default agreement's effect on competition depend partly on which platform obtains default status. When it is the dominant platform, the agreement magnifies its perceived quality advantage over laggard competitors. This allows the dominant platform to increase unwanted advertising activities without losing too many consumers, analogous to a monopolist charging a large markup. By contrast, when a laggard acquires default status, the dominant platform's quality advantage shrinks, forcing it to behave more competitively. However, the default contract achieves this by effectively reducing the dominant platform's quality, so consumers do not necessarily benefit.

Another relevant factor is the breadth of a platform's default agreements—the share of search traffic that they cover. In the baseline model (Sections 2-2.4), I consider defaults that are as broad as possible in the sense that they apply to all search queries, regardless of what browser or other product a consumer uses to access her preferred search engine. Under these circumstances, a default agreement involving a laggard platform will always reduce consumer welfare. In fact, for most parameter values, the laggard's contract is worse for consumers than an identical agreement favoring the dominant platform. One key reason for this is simply that the dominant platform is more popular.[3] Thus, a broad set of defaults involving a dominant platform does not necessarily leave consumers worse off if the counterfactual alternative would involve a laggard acquiring default status on the same broad scale.

In Section 3, I consider narrower default contracts that apply to a single browser. This leads to spillover effects, as a default contract on one browser will generally affect users on all other browsers. Interestingly, even if a default contract on a given browser harms all consumers who use that browser (which is typically true), it may have positive spillovers on users of other browsers. In fact, a laggard's default contract will always generate positive spillovers by forcing the dominant platform to behave more competitively on all browsers. As a result, the laggard's default contract can raise aggregate consumer welfare if it is sufficiently narrow. By contrast, the dominant platform's default contract will still generally harm consumers (even if it is narrow), provided that algorithmic learning effects are not too strong.

Algorithmic learning magnifies the competitive effects of default agreements, because it makes consumer demand much more responsive to small changes in most relevant variables. For example, a modest asymmetry in platform quality could produce a large imbalance in market shares. Consequently, a default agreement can have significant competitive effects even if

[3] This means that, when the dominant platform becomes the default, most consumers are nudged toward the same option they would have picked anyway. Hence, most consumers are not directly harmed by switching costs, although they are still typically harmed by a reduction in competition.

switching costs are relatively low. If algorithmic learning effects are sufficiently strong, then consumers could be better off under monopoly than competition.[4] In this case, the dominant platform's default agreements will benefit consumers precisely because they reduce competition. However, most research suggests that algorithmic learning is subject to diminishing marginal returns. In an extension, I show that this makes it more likely that the dominant platform's default agreements reduce consumer welfare (Section 5.1).

After exploring how defaults affect competition and consumer welfare, I show that advertisers usually benefit from default agreements involving the dominant search platform, but not those involving a laggard. Competition for consumers leads platforms to cut back on advertising activities that consumers dislike. While good for consumers, this makes the platform's advertising services less valuable to advertisers. Hence, advertisers tend to benefit from agreements that soften competition on the consumer side.

I also consider dynamic incentives for entry and investment (Section 4). Default contracts struck by laggard platforms can make entry easier and may therefore stimulate competition in the long run. This is consistent with empirical work finding that vertical restraints by laggards can help to facilitate competitive entry in network industries (e.g. Lee, 2013). By contrast, the dominant platform's default agreements make it harder for laggards to enter. They also reduce laggards' incentive to invest in quality. However, the dominant platform's agreements can potentially help to facilitate entry in adjacent markets. I also consider the possibility that Google's large default payments could discourage device makers (e.g. Apple) from vertically integrating into the general search market.

Section 6 discusses policy implications. One key question is of what remedy or regulatory intervention (if any) would be most appropriate. I discuss several possibilities that have been proposed. A categorial ban on default agreements—or a policy mandating choice screens on all devices and browsers—may not be the best solution. This is because, in small doses, default status (particularly for laggards) may be a competitive stimulant, as noted above. This section also discusses why search engine defaults are unlikely to result in lower prices of mobile devices.[5] Finally, I emphasize that this paper does not suggest that default-like arrangements in other industries are likely to raise antitrust concerns.

Related Literature

There is relatively little economic scholarship on search engine defaults. Decarolis et al. (2024) studies Google's recent implementation of a choice screen on Android devices in the European Economic Area (EEA), Russia, and Turkey. They find that this caused a drop in Google's market share, which ranges from less than 2 percent to as much as 12 percent, depending on the country and the choice screen. They also find no evidence that choice screen led to an increase in the price of devices. The EEA choice screen was initially organized as an auction in which search engines pay to appear on the screen. Ostrovsky (2023) presented a model of that

[4] This is a possibility that arises in many models of competition in network industries. (see, e.g., Nocke et al., 2007; Weyl and White, 2014).

[5] I summarize the main reasons for this in Section 6. I provide a more detailed discussion in Appendix B.3.

auction. However, Google has since eliminated the auction and made it free to competitors.

In an antitrust policy article, Bet et al. (2022) argue that search engine defaults could harm consumers by degrading platform quality and undermining incentives for innovation. Choi et al. (2023) provide a model of a tying arrangement involving a monopoly good and a tied product that is subject to network effects. This model sheds light on Google's requirement that device makers preinstall a Google Search app as a condition for gaining access to the Google Play store. The European Commission's investigation into this app store bundling led to an order forcing Google to implement a choice screen on Android devices in the EEA.

The most similar paper to this one is likely Chen and Schwartz (2024), which was posted online a few weeks before this paper. They also provide a model of default agreements that nudge consumers toward a particular brand. But there are substantial differences between our papers. They do not focus specifically on search engines, so most of their paper does not consider network effects (i.e. algorithmic learning), which is a core issue in my paper. They consider homogeneous one-sided firms, whereas I focus on differentiated two-sided platforms. And they focus heavily on the process by which firms bid for default status, while I do not. Overall, our papers are complementary but distinct.

A number of theory papers focus on other aspects of search engine markets, such as auctions for ad slots on a search page (e.g. Athey and Ellison, 2011) Other articles consider how search engines may strategically distort or "obfuscate" search results to make more money from advertisers (Garcia, 2023). More generally, some papers consider how digital platforms may compete for consumer attention (e.g. Prat and Valletti, 2022).

At a broad level, this paper contributes to the literature on two-sided media platforms. The literature is surveyed by Anderson and Jullien (2015) and Peitz and Reisinger (2015).

2 Baseline Model

There are two search platforms $i = 1, 2$, which are vertically and horizontally differentiated. Platform 1 is dominant. Platform 2 is a laggard. The market is two-sided, with a unit mass of consumers on one side and a mass $m > 0$ of advertisers on the other. Algorithmic learning effectively creates direct (within-group) network effects on the consumer side. The model is of the well-known "competitive bottleneck" format,[6] meaning that consumers singlehome (they participate on a single platform), whereas advertisers can multihome. This is a good fit for the present context, as most consumers use a single search engine but advertisers often run ads on more than one platform.

The "baseline model" presented in this section assumes that a default agreement is maximally broad—it affects all consumers equally, regardless of what browser they use. Hence, at most one platform can have default status. This simplification is useful for studying the case where a single dominant platform acquires the default position on almost all major browsers. The

[6]See Armstrong (2006); Armstrong & Wright (2007).

subsequent section explores the more general case where a default contract applies to a single browser, and allows for the possibility that both platforms have defaults (on different browsers).

The absence of consumer-side prices raises the question of how exactly platforms compete for consumers. For example, on the consumer side of the market, what is the analogue to a price increase? The literature offers a number of possible answers:

- The platform could make the ads on its search page more numerous, more prominent, or harder to ignore. This benefits advertisers, but it subjects consumers to "nuisance costs" (e.g. Anderson and Jullien, 2015) or "attention costs" (e.g. Newman, 2015; Gal and Rubinfeld, 2016).
- The platform could gather more data on its users. This is good for advertising but bad for consumers who value their privacy (e.g. Chen and Schwartz, 2024).
- The platform could strategically distort its search results (at the expense of accuracy or relevance) to induce more engagement with ads (e.g. Garcia, 2023) or to enable advertising sellers to charge higher prices (e.g. De Corniere, 2016).

At a broad level, all of these behaviors generate the same basic tradeoff: they enhance the platform's advertising services, but they degrade the consumer's experience. I attempt to generalize such behaviors by introducing a strategic choice variable—the *ad intensity level*, denoted α_i for each platform i—that generates precisely this tradeoff. In effect, α_i is a quality variable that affects the two sides in opposite ways: an increase in α_i makes the platform *more* attractive to advertisers, but *less* attractive to consumers. In other words, α_i governs the extent to which the platform prioritizes advertising over consumer enjoyment.

For example, following the examples listed above, α_i could represent the number of ads on the search page; the amount of personal data that the platform gathers on search users; or the extent to which search results are distorted in ways that benefit advertisers.

The rest of this section lays out the baseline model in detail. The following table will help to keep track of key notations.

Notation	Definition
x	A consumer's type (her location on a Hotelling line)
z	An advertiser's type
X_i	Consumer demand for platform i
Z_i	Advertiser demand for platform i
$F(z), f(z)$	The CDF and PDF of z
p_i	The price of advertising services on platform i
α_i	The ad intensity level on platform i
$\tilde{z}_i \equiv \alpha_i - p_i$	The marginal advertiser type on platform i
V_i	The exogenous quality of platform i's search technology
σ	The cost of switching from the default platform to the non-default platform
δ_i	A dummy with $\delta_i = 1$ if platform i is the default and $\delta_i = 0$ otherwise
Δ	Platform 1's technological advantage over platform 2
η	The strength of network effects
$\mu(\tilde{z}_i) \equiv F(\tilde{z}_i)/f(\tilde{z}_i)$	The inverse semi-elasticity of advertising demand on platform i

2.1 Advertisers

An advertiser has a type $z \in \mathbb{R}$, which is distributed according to CDF $F(z)$. We assume F is twice continuously differentiable with density denoted $f = \partial F/\partial z$. The support of F is assumed to be an interval with a finite lower bound, denoted $\underline{z}$. It may or may not have a finite upper bound. Thus, the type space could be compact, $[\underline{z}, \overline{z}]$, or infinite, $[\underline{z}, \infty)$.

Advertisers with higher types benefit less from advertising, and hence have a lower willingness to pay. Specifically, the gross benefit that an advertiser gets from an ad campaign on platform i is $(\gamma + \alpha_i - z)X_i$. Here $X_i \in [0, 1]$ is consumer demand for platform i, which specifies the measure of consumers who participate on the platform. The parameter $\gamma \geq 0$ captures any benefits of advertising that are unrelated to the ad intensity level (α_i). However, because we allow $\underline{z}$ to be negative, we can normalize $\gamma = 0$ without loss of generality.[7] The total price of running an ad campaign on platform i is $X_i p_i$, where p_i is the price paid for each consumer on the platform. Thus, the net payoff an advertiser gets from joining platform i is

$$\pi_i(z) = (\alpha_i - z - p_i)X_i. \tag{1}$$

The total payoff an advertiser gets from multihoming is simply the sum $\pi_1(z) + \pi_2(z)$. Thus, an advertiser's participation decisions are independent: it will participate on platform i if and only if $\pi_i(z) \geq 0$, regardless of whether it also joins platform j. The measure of advertisers who join platform i, denoted Z_i, is therefore

$$Z_i \equiv m \Pr\Big(\pi_i(z) \geq 0\Big) = mF(\alpha_i - p_i). \tag{2}$$

This is the demand function for advertising services on platform i. Note that in general the

[7]We could always redefine advertiser types as $\hat{z} = z - \gamma$.

advertiser side of the market is not fully covered (some advertisers do not join either platform).

2.2 Consumers and Defaults

Consumers are uniformly distributed along a Hotelling line, with locations indexed by $x \in [0, 1]$. Platform 1 (resp. platform 2) is located at $x = 0$ (resp. $x = 1$). Travel costs take the linear form $t \times distance$, where $t > 0$. A platform's overall perceived quality among consumers has three components:

$$\text{Quality}_i = V_i + X_i\eta - h\alpha_i \tag{3}$$

The term V_i is the exogenous component of platform quality. This represents the quality of platform i's search technology, such as the design of its algorithm. The second term, $X_i\eta$, captures network effects, with $\eta \geq 0$ giving the rate at which quality rises as a platform gains more users. Finally, $h\alpha_i$ (where h is a positive scalar) captures the disutility consumers experience due to the ad intensity level α_i. Intuitively, one can think of the disutility $h\alpha_i$ as the "price" a consumer pays to use platform i's search engine.

Because platform 1 is dominant, we assume $V_1 > V_2$. This will ensure that it attracts a majority share of consumers in equilibrium ($X_1 > X_2$). We also assume the consumer side of the market is fully covered in all equilibria: $X_1 + X_2 = 1$.

When one platform is the default, a consumer who wishes to use the other platform must incur a switching cost of $\sigma \geq 0$. Let $\delta_i = 1$ if platform i is the default and $\delta_i = 0$ otherwise. Both platforms cannot be the default, and hence $\delta_1 + \delta_2 \leq 1$. But we allow for the possibility that neither platform has default status: $\delta_1 = \delta_2 = 0$.

A consumer at location x gets a utility of $u_i(x)$ from joining platform i, where

$$u_1(x) = V_1 - \delta_2\sigma + X_1\eta - h\alpha_1 - tx, \qquad u_2(x) = V_2 - \delta_1\sigma + X_2\eta - h\alpha_2 - t(1 - x). \tag{4}$$

The marginal consumer, defined by $u_1(\widetilde{x}) = u_2(\widetilde{x})$, is located at

$$\widetilde{x} = \frac{V_1 - \delta_2\sigma - h\alpha_1 - (V_2 - \delta_1\sigma - h\alpha_2) + 2X_1\eta + t - \eta}{2t}, \tag{5}$$

Consumer demand levels are pinned down as $X_1 = \widetilde{x}$ and $X_2 = 1 - \widetilde{x}$. This gives rise to the following expressions for consumer demand:

$$X_1 = \frac{1}{2} + \frac{\Delta - h(\alpha_1 - \alpha_2)}{2(t - \eta)} \qquad X_2 = \frac{1}{2} - \frac{\Delta - h(\alpha_1 - \alpha_2)}{2(t - \eta)}, \tag{6}$$

where Δ is defined by

$$\Delta \equiv V_1 - V_2 + (\delta_1 - \delta_2)\sigma. \tag{7}$$

Here Δ is the dominant platform's *technological advantage* over the laggard. This definition is useful because we can think of a default as effectively reducing one of the quality levels V_1 or V_2. Specifically, when platform i acquires default status, this effectively reduces V_j by σ. Thus, a default either increases or decreases the dominant platform's technology advantage, depending on which platform is the default.

By inspection of (6), the biasing effects of defaults manifest as vertical shifts in the demand functions X_i. The default (resp. non-default) platform's demand function shifts upward (downward). The magnitude of these demand shifts is $\sigma/2(t-\eta)$. Notice that, if algorithmic learning effects are strong (large η), then these shifts can be large even if the switching costs σ is small.

Throughout the paper we assume that $t > \eta$. This is necessary for both platforms to be active in equilibrium.[8] For simplicity, we will also assume that σ is not large enough to make Δ negative when the laggard obtains default status (i.e. we assume $\sigma < V_1 - V_2$). As noted above, we allow for the possibility that neither platform is preselected by default ($\delta_1 = \delta_2 = 0$). We will associate this possibility with a choice screen. A choice screen treats different search platforms the same, and hence does not generate a biasing effect. An example of a choice screen is depicted in Appendix B.1.

2.3 Platform Profits and Competition

Platforms earn all revenues from advertisers, since they do not charge prices to consumers. For simplicity, we normalize platforms' production costs to zero. Platforms compete by simultaneously choosing p_i and α_i. However, it is equivalent, and more convenient, to think of platform i as choosing p_i and $\tilde{z}_i$, where $\tilde{z}_i \equiv \alpha_i - p_i$ is the marginal advertiser type on platform i (i.e. the type such that $\pi_i(\tilde{z}_i) = 0$). Notice that $\tilde{z}_i$ is a good proxy for advertiser demand (Z_i), because they are directly related by $Z_i = mF(\tilde{z}_i)$. Once $\tilde{z}_i$ and p_i are chosen, the ad intensity level is pinned down as $\alpha_i = \tilde{z}_i + p_i$.

Platform i's profits are given by

$$\Pi_i = Z_i X_i p_i = mF(\tilde{z}_i) X_i p_i, \tag{8}$$

where consumer demand is now expressed in terms of $\tilde{z}_i$ and p_i as

$$X_1 = \frac{1}{2} + \frac{\Delta - h(\tilde{z}_1 + p_1 - \tilde{z}_2 - p_2)}{2(t-\eta)}, \qquad X_2 = \frac{1}{2} - \frac{\Delta - h(\tilde{z}_1 + p_1 - \tilde{z}_2 - p_2)}{2(t-\eta)}. \tag{9}$$

In what follows, we will make frequent use of the following function:

$$\mu(z) \equiv \frac{F(z)}{f(z)}. \tag{10}$$

Evaluated at $\tilde{z}_i$, this gives the inverse semi-elasticity of demand for advertising services on platform i.[9] As others have noted, the function $\mu(\cdot)$ is useful in characterizing demand (e.g. Johnson, 2017). Also, its equilibrium value is a good proxy for market power (e.g. Weyl, 2010).

[8] If $\eta \geq t$, then network effects are so strong that the market always "tips" into monopoly.

[9] To see this, note that

$$\mu(\tilde{z}_i) = \frac{F(\alpha_i - p_i)}{f(\alpha_i - p_i)} = \frac{mF(\alpha_i - p_i)}{mf(\alpha_i - p_i)} = -\frac{Z_i}{\frac{\partial Z_i}{\partial p_i}} = \frac{p_i}{\varepsilon_i},$$

We close this section with three technical assumptions on μ. First, we assume that advertising demand (Z_i) is strictly log-concave, which is equivalent to the condition that $\mu'(\cdot) > 0$.

Assumption 1: $\mu(z)$ is strictly increasing.

The derivative $\mu'(z)$ could be increasing or decreasing (or nonmonotonic). Hence, $\mu(z)$ could be concave or convex. However, our second assumption stipulates that it cannot be *too* convex.

Assumption 2: $\dfrac{\mu'(z)}{\mu(z)}$ is nonincreasing.

This says that $\mu(z)$ is (weakly) log-concave. Both assumptions 1 and 2 are satisfied by most well-known distributions with support bounded from below.[10] In Section 2.4 we will give an economic justification for assumption 2 based on revealed preference (see footnote 22). Our final assumption simply ensures that each platform's optimal choice of $\tilde{z}_i$ is always interior.

Assumption 3: $F(\tilde{z}^m) < 1$, where $\tilde{z}^m \equiv \mu^{-1}\left(\dfrac{2(t-\eta)}{h}\right)$.

2.4 Results

This section presents the main results for the baseline model. We focus on equilibria in which both platforms are active, which requires that Δ is not too large.

Proposition 1. *There exists a critical threshold $\Delta^{crit} > 0$ such that the game has a unique equilibrium with both platforms active whenever $\Delta < \Delta^{crit}$. This equilibrium is characterized by the equations*

$$p_i^* = \mu(\tilde{z}_i^*) \tag{11}$$

and

$$X_i^* = \frac{h}{2(t-\eta)}\mu(\tilde{z}_i^*). \tag{12}$$

Equation (11) is simply the standard inverse elasticity pricing rule.[11] Equation (12) is more interesting—it specifies how market outcomes are correlated across the two sides. Plugging the first equation into the second, it says that consumer-side market shares (X_i^*) are proportional to advertising prices (p_i^*). Given that the consumer side of the market is fully covered, this leads to the following result:

$$p_1^* + p_2^* = \frac{2(t-\eta)}{h}. \tag{13}$$

Thus, the sum of the platforms' ad prices is independent of Δ, and hence is unaffected by default agreements. By contrast, the individual prices do depend on Δ. In particular, the price

where $\varepsilon_i \equiv -\left(\frac{\partial Z_i}{\partial p_i}\right)\frac{p_i}{Z_i}$ is the price elasticity of advertising demand on platform i.

[10] Examples include the uniform, exponential, lognormal, Pareto, and chi-squared distributions, among others.

[11] Recall that $\mu(\tilde{z}_i)$ is equal to platform i's inverse semi-elasticity, i.e. $\mu(\tilde{z}_i) = p_i/\varepsilon_i$. Then (11) can be written as $\varepsilon_i^* = 1$, which is simply the inverse elasticity pricing rule in the special case of costless production.

ratio is equal to the ratio of consumer-side market shares:

$$\frac{p_1^*}{p_2^*} = \frac{X_1^*}{X_2^*}. \tag{14}$$

As an illustration, consider the equilibrium in the case where the model is linear.

Example (linear demand). *Suppose z is uniformly distributed over $[0, \ell]$, so that $F(z) = z/\ell$. Then the equilibrium strategies and demand levels are*

$$\tilde{z}_1^* = p_1^* = \frac{t-\eta}{h} + \frac{\Delta}{5h} \qquad \tilde{z}_2^* = p_2^* = \frac{t-\eta}{h} - \frac{\Delta}{5h}$$

$$\alpha_1^* = 2\left(\frac{t-\eta}{h} + \frac{\Delta}{5h}\right) \qquad \alpha_2^* = 2\left(\frac{t-\eta}{h} - \frac{\Delta}{5h}\right)$$

$$X_1^* = \frac{1}{2} + \frac{\Delta}{10(t-\eta)} \qquad X_2^* = \frac{1}{2} - \frac{\Delta}{10(t-\eta)}$$

$$Z_1^* = \frac{m}{\ell}\left(\frac{t-\eta}{h} + \frac{\Delta}{5h}\right) \qquad Z_2^* = \frac{m}{\ell}\left(\frac{t-\eta}{h} - \frac{\Delta}{5h}\right)$$

The fact that $\tilde{z}_i^* = p_i^*$ in this example reflects the fact that, when the model is linear, $\mu(\cdot)$ is the identify function.

2.4.1 Comparative Statics

When $\Delta = 0$, the equilibrium is fully symmetric.[12] The proposition below clarifies how changes in Δ—and, by extension, default contracts—affect the equilibrium. It also explains how the strength of algorithmic learning (η) affects the equilibrium.

Proposition 2. *The equilibrium exhibits the following comparative statics:*

(i) *$\tilde{z}_1^*$, p_1^*, α_1^*, and X_1^* are strictly increasing in Δ, while $\tilde{z}_2^*$, p_2^*, α_2^*, and X_2^* are strictly decreasing in Δ. The share-weighted average level of ad intensity, $X_1^*\alpha_1^* + X_2^*\alpha_2^*$, is also strictly increasing in Δ.*

(ii) *$\tilde{z}_i^*$, p_i^*, α_i^* $(i = 1, 2)$, and X_2^* are strictly decreasing in η, while X_1^* is strictly increasing in η.*

The comparative statics in part (i) tell us how default agreements affect market outcomes. Giving default status to the dominant platform increases its market power, allowing it to raise its ad price and its ad intensity level while simultaneously capturing more users on both sides. By contrast, the laggard gets weaker—it loses users on both sides, and lowers both its ad intensity level and ad price. These effect imply that the consumer side of the market becomes

[12] See the proof of Proposition 1.

more concentrated. By contrast, when the laggard acquires default status, all effects just mentioned are reversed.

Part (i) also implies that competition for consumers grows softer when the dominant platform acquires default status. The increase in the share-weighted average ad intensity level means that consumers are subjected to more unwanted advertising activity on average. This is analogous to an increase in the share-weighted average price level in a traditional market. By contrast, when the laggard becomes the default, competition for consumers becomes more intense. However, it is important to keep in mind that a default's impact on competitiveness is just one of several channels through which it affects consumer welfare. We consider all the different channels in the next section.

Because part (i) implies $\tilde{z}_1^* > \tilde{z}_2^*$, it follows that all advertisers who join platform 2 are multihomers. All other advertisers either singlehome on platform 1 (if $\tilde{z}_2^* < z \leq \tilde{z}_1^*$) or else they are inactive (if $z > \tilde{z}_1^*$). Therefore, when the dominant platform acquires default status, this increases the total number of advertisers who are active (by increasing $\tilde{z}_1^*$), but it also reduces multihoming (by decreasing $\tilde{z}_2^*$).

Part (ii) of Proposition 2 can be understood as follows. When η increases, platforms compete more fiercely for consumers, leading both to reduce their ad intensity levels, α_i^*. From advertisers' perspective, this is like a reduction in the quality of advertising services. This shifts the advertising demand functions vertically downward, inducing the platforms to cut ad prices, albeit not by enough to avoid a drop in ad sales. By inspection of (9) increasing η makes consumer demand X_i more sensitive to the dominant platform's technology advantage, Δ. Thus, a given change in Δ will lead more consumers to switch platforms when η is larger.[13] In other words, algorithmic learning magnifies a default's impact on consumer-side market shares. This can be understood in terms of a feedback loop.[14]

This has an important policy implication: due to algorithmic learning, a default agreement can have significant competitive effects even if switching costs are not very large.[15] As a result, these deals may raise antitrust concerns even though they don't restrain consumer choice as forcefully as some more familiar vertical restraints (e.g. exclusive dealing).

2.4.2 Consumer Welfare

Aggregate consumer welfare, denoted $\mathcal{W}_c$, is

$$\mathcal{W}_c = \int_0^{X_1^*} u_1^*(x)dx + \int_{X_1^*}^1 u_2^*(x)dx, \tag{15}$$

[13]One consequence of this is that the maximal sustainable quality gap Δ^{crit} from Proposition 1 is strictly decreasing in η.

[14]When platform i becomes the default, some consumers switch from platform j to platform i. Due to algorithmic learning, this increases i's quality and reduces j's quality. These quality shifts then lead even more consumers to switch from j to i, which generates additional quality shifts, and so on ad infinitum.

[15]This was foreshadowed by our earlier observation that a default generates vertical shifts of $\sigma/2(t-\eta)$ in the consumer-side demand functions (Section 2.2). This shows that σ and η contribute independently to a default's impact on consumer demand.

where $u_1^*(x)$ and $u_2^*(x)$ are the equilibrium utilities a consumer at location x would get for joining platform 1 or 2, respectively.[16] We want to know how a default contract affects consumer welfare relative to the benchmark where neither platform is the default (e.g. because consumers are presented with a choice screen). The answer is not obvious, as there are three distinct channels through which defaults affect consumer utility: competition effects (changes in ad intensity), algorithmic learning effects, and switching costs.

Suppose platform i is the default ($\delta_i = 1$). Consider the marginal effect of a small increase in the switching cost σ. It is easy to verify that[17]

$$\frac{\partial \mathcal{W}_c}{\partial \sigma} = X_i^* \frac{\partial u_i^*}{\partial \sigma} + X_j^* \frac{\partial u_j^*}{\partial \sigma}. \tag{16}$$

Hence, the marginal effect on aggregate consumer welfare is just the share-weighted sum of the utility effects experienced by users on each platform.[18] Consider first the marginal utility effect on users of the default platform:

$$\frac{\partial u_i^*}{\partial \sigma} = \underbrace{\frac{\partial X_i^*}{\partial \sigma}\eta}_{\substack{\text{learning} \\ \text{effect } (>0)}} \quad \underbrace{-\,h\frac{\partial \alpha_i^*}{\partial \sigma}}_{\substack{\text{competition} \\ \text{effect } (<0)}} \tag{17}$$

These users experience a positive algorithmic learning effect, since the default platform gains market share (X_i^* increases). But they also experience a negative competition effect, because the default platform gains market power, leading it to raise ad intensity (α_i^* increases). Next, consider the marginal utility effect on users of the non-default platform, j:

$$\frac{\partial u_j^*}{\partial \sigma} = \underbrace{-1}_{\substack{\text{switching} \\ \text{cost effect} \\ (<0)}} \quad \underbrace{+\frac{\partial X_j^*}{\partial \sigma}\eta}_{\substack{\text{learning} \\ \text{effect } (<0)}} \quad \underbrace{-\,h\frac{\partial \alpha_j^*}{\partial \sigma}}_{\substack{\text{competition} \\ \text{effect } (>0)}} \tag{18}$$

Thus, users on the non-default platform also experience a learning effect and a competition effect, but the sign of each effect is now reversed, since the default leads X_j^* and α_j^* to fall. These users also experience a third marginal effect: because there is a cost σ of adopting the non-default platform, users on that platform sustain a negative switching cost effect, $\frac{\partial(-\sigma)}{\partial \sigma} = -1$.

Proposition 3. *Suppose platform i is the default, and let $j \neq i$.*

(i) $\dfrac{\partial u_j^*}{\partial \sigma} < 0$ *in all cases.*

(ii) If $\eta \leq \frac{2}{3}t$*, then* $\dfrac{\partial u_i^*}{\partial \sigma} < 0$*. But if* η *is sufficiently larger than* $\frac{2}{3}t$*, then* $\dfrac{\partial u_i^*}{\partial \sigma} > 0$*.*

Thus, unsurprisingly, a default always harms consumers on the non-default platform.[19] As for

[16] Explicitly: $u_1^*(x) = V_1 - \delta_2\sigma + X_1^*\eta - h\alpha_1^* - tx$ and $u_2^*(x) = V_2 - \delta_1\sigma + X_2^*\eta - h\alpha_2^* - t(1-x)$.

[17] This uses the fact that $u_1^*(X_1^*) = u_2^*(X_1^*)$.

[18] Note that the marginal utility effects $\frac{\partial u_k^*}{\partial \sigma}$ do not depend on the consumer's location x.

[19] In fact, for these users, the negative switching cost effect always dominates the positive competition effect,

users of the default platform, their marginal utiliy effect can be positive, but only if η is quite large. The value of η at which $\partial u_i^*/\partial\sigma$ changes sign can fluctuate as σ grows, but it is always strictly greater than $\frac{2}{3}t$ and strictly less than t. (See the proof for details.) We can use these results to determine the cumulative effect of the default on consumer welfare.

Proposition 4. *Defaults affect aggregate consumer welfare ($\mathcal{W}_c$) as follows:*

(i) If the laggard is the default, then $\mathcal{W}_c$ strictly falls in all cases.

(ii) If the dominant platform is the default, then $\mathcal{W}_c$ strictly falls whenever $\eta \leq \frac{2}{3}t$. But if η is sufficiently larger than $\frac{2}{3}t$, then $\mathcal{W}_c$ strictly increases.

For the dominant platform's default to increase consumer welfare, η must be very large. The gains to users on the dominant platform must outweigh the harm to other consumers. However, the fact that welfare rises when η is very large has little to do with defaults specifically. Rather, it reflects a property of the market itself: when algorithmic learning effects are very strong, we are effectively in a natural monopoly regime where consumers are better off under monopoly than under competition. In this regime, *any* anticompetitive arrangement that increases market concentration (e.g. a merger) will raise consumer welfare. An additional caveat is that, if algorithmic learning benefits diminish with scale (as most research suggests), it becomes much less likely that the dominant platform's default could raise consumer welfare (Section 5.1).

In the baseline game, the laggard's default can never raise consumer welfare. For welfare to rise, the default must generate a large increase in algorithmic learning benefits. If we add up these benefits for all consumers in the market, we get $X_1^*(X_1^*\eta) + X_2^*(X_2^*\eta) = HHI \times \eta$, where HHI is the Herfindahl index. Thus, for a default to increase consumer welfare, it must increase market concentration. But a default by the laggard does just the opposite, since it makes the market more competitive. Importantly, however, these are static results—they take the laggard's existence as given. As we show later, when entry is endogenous, a laggard's default contract can benefit consumers by facilitating entry (Section 4). Additionally, when defaults have limited scope, a laggard's default can increase static consumer welfare (Section 3).

Proposition 4 implies that, when η is very large, consumers prefer a default favoring the default platform to a default favoring the laggard. But what if η is not that large, so that a default by either platform would reduce consumer welfare? Even in this case, consumers will *usually* prefer the dominant platform to be the default, rather than the laggard. This is because the dominant platform is more popular ($X_1^* > X_2^*$). Hence, when it acquires default status, most consumers are nudged toward the same option they would have picked anyway.[20] Despite this, it is possible for the dominant platform's default to be more harmful than the laggard's, although this requires a rather demanding condition.

Proposition 5. *If η is sufficiently small and $\mu''(\cdot)$ is sufficiently negative, then a default favoring the dominant platform is more harmful to consumers than a default favoring the laggard.*

so their utility would fall even without a negative learning effect.

[20]Thus, most consumers avoid the switching cost. But they are still harmed by a reduction in competition.

Our assumptions do not constrain the sign of $\mu''(\cdot)$, so most allowed specifications of $\mu(\cdot)$ won't satisfy the condition in Proposition 5. The intuition for the result is that, if $\mu(\cdot)$ is highly concave ($\mu''(\cdot) \ll 0$), then a change in Δ elicits a much larger change in α_1^* than in α_2^*, which makes the dominant platform's default more painful and the laggard's less painful.[21]

2.4.3 Advertiser Surplus

How do default agreements affect advertisers? Aggregate advertiser surplus, denoted $\mathcal{W}_a$, is

$$\mathcal{W}_a = \int_{\underline{z}}^{\widetilde{z}_1^*} \pi_1^*(z) f(z) dz + \int_{\underline{z}}^{\widetilde{z}_2^*} \pi_2^*(z) f(z) dz. \tag{19}$$

Here $\pi_i^*(z)$ is the equilibrium payoff a type-z advertiser gets from joining platform i, while $\underline{z}$ is the lower bound on the support of F. It turns out that advertisers like default agreements when they favor the dominant platform, but not if they favor the laggard.

Proposition 6. *A default favoring the dominant platform strictly increases aggregate advertiser surplus ($\mathcal{W}_a$), whereas a default favoring the laggard strictly reduces it.*

This result reflects two things. First, recall that a default for the dominant platform (resp. laggard) makes competition for consumers less (more) intense. But advertisers don't like it when consumer-side competition is intense, since this puts downward pressure on ad intensity levels (α_i), which feels like a quality reduction to advertisers. Second, competition on the advertiser side is always soft, even if competition on consumer side is intense. This reflects a well-known feature of platform competition models where one side (consumers) singlehomes and the other side (advertisers) can multihome: the multihomers do not view the platforms as close substitutes, since each confers access to a distinct set of singlehomers to interact with (Armstrong, 2006; Armstrong and Wright, 2007). As a result, the usual harms associated with reduced competition show up mainly on the singlehoming side.

2.4.4 Platform Willingness To Pay for Defaults

Both platforms would benefit from being the default and hence would be willing to pay for that privilege. However, the dominant platform is willing to outbid the laggard.[22]

Proposition 7. *Platform 1's willingness to pay for default status is strictly larger than that of platform 2.*

[21] This is explained in detail in Appendix B.2.

[22] The proof hinges on the fact that μ is log-concave (Assumption 2). If that assumption were violated, μ could be extremely convex to the point that a small increase in Δ generates virtually zero increase in α_1^*, but a very large reduction in α_2^*. (For an explanation of this, see Appendix B.2.) In that case, the increase in Δ could reduce the laggard's profits by more than it raises the dominant platform's profits. We would then expect the laggard to outbid the dominant platform. However, in practice, Google outbids all smaller search engines for default status on most major web browsers. This provides some additional justification for Assumption 2.

This result can be viewed as an application of Gilbert and Newbery (1982), which explained that a dominant firm generally gains more profits by maintaining its technological advantage over a laggard than the laggard would gain by catching up to the dominant firms.

3 Browser-Specific Defaults and Spillovers

The baseline model considers defaults that are as broad as possible—they affect all consumers equally, no matter what browser they use. This is useful for considering the case where a single platform acquires the default position on most or all browsers. But in practice, a platform's default coverage need not be that broad. To that end, we now generalize the model so that a default contract applies to a single browser, and thus directly affects only subset of consumers. We also allow for the possibility that both platforms have defaults (on different browsers).

As demonstrated below, this leads to some qualitatively new effects. First, in this setting, default agreements generate spillovers—a default on one browser affects users on other browsers. And, interestingly, these spillovers can be positive—and can potentially lead aggregate consumer welfare to rise—even if the default harms all consumers on the browser it applies to. Second, this generalization allows us to consider how the *breadth* of a default agreement (the share of consumers directly affected by it) shapes its effects. Finally, it shows how competing defaults "interact" in the sense that each one influences the welfare effects of the other.

We now suppose there are two search access points, which we will call "browser 1" and "browser 2." The fraction of consumers who use browser k is $\lambda^k \geq 0$, where $\lambda^1 + \lambda^2 = 1$. We are interested in the case where each search engine has default status on at most one browser.[23] To this end, we will assume that platform 1 is the default on browser 1, while platform 2 is the default on browser 2. The switching costs from these deals are denoted $\sigma^1 \geq 0$ and $\sigma^2 \geq 0$, respectively.[24] We make the following assumptions:

- A consumer's browser selection is exogenously fixed and independent of her location, x.
- Each search platform chooses a single ad intensity level α_i, which applies to both browsers.
- Network benefits depend on the total number of consumers who use a given search platform, not on how they are allocated between browsers.
- An advertiser's payoff from using a given platform depends on the total number of consumers who use that platform, not on how they are allocated between browsers.[25] Consequently, advertiser payoffs are the same as in the baseline game.

Consider a consumer at location x who uses browser k. Her utility from joining search platform

[23] The baseline game already studied the case where one platform is the default on all access points.

[24] In general $\sigma^1 \neq \sigma^2$, as it may be easier to switch search engines on one browser than the other. Also, as noted above, a default contract may include restrictions that raise switching costs (footnote 2).

[25] In other words, a search platform gives advertisers equal access to all of its consumers across all browsers.

i is denoted $u_i^k(x)$. Thus, there are four utility functions, one for each browser-search platform combination. They are:

$$u_1^1(x) = V_1 + \eta X_1 - h\alpha_1 - tx \qquad u_2^1(x) = V_2 - \sigma^1 + \eta X_2 - h\alpha_2 - t(1-x)$$
$$u_1^2(x) = V_1 - \sigma^2 + \eta X_1 - h\alpha_1 - tx \qquad u_2^2(x) = V_2 + \eta X_2 - h\alpha_2 - t(1-x)$$

To make sense of this, recall that platform i is the default platform on browser k if and only if $i = k$. Thus, whenever $i \neq k$, the utility u_i^k includes a switching cost $-\sigma^k$. As in the baseline game, X_i represents the total number of consumers who use search platform i. The value of X_i can be decomposed as

$$X_i = \lambda^1 X_i^1 + \lambda^2 X_i^2, \tag{20}$$

where X_i^k is platform i's market share on browser k. These shares are given by $X_1^k \equiv \widetilde{x}^k$ and $X_2^k \equiv 1 - \widetilde{x}^k$, where $\widetilde{x}^k$ is the marginal consumer type on browser k.[26] It is easy to verify that

$$\widetilde{x}^k = \frac{\Delta^k + 2\eta X_1 + t - \eta - h(\alpha_1 - \alpha_2)}{2t}, \tag{21}$$

where $\Delta^1 \equiv V_1 - V_2 + \sigma^1$ and $\Delta^2 \equiv V_1 - V_2 - \sigma^2$. This leads to the following expressions for aggregate demand:

$$X_1 = \frac{1}{2} + \frac{\lambda^1\Delta^1 + \lambda^2\Delta^2 - h(\alpha_1 - \alpha_2)}{2(t-\eta)}, \quad X_2 = \frac{1}{2} - \frac{\lambda^1\Delta^1 + \lambda^2\Delta^2 - h(\alpha_1 - \alpha_2)}{2(t-\eta)}. \tag{22}$$

These are very similar to the consumer demand functions from the baseline game. The only difference is that Δ has been replaced by the convex combination $\lambda^1\Delta^1 + \lambda^2\Delta^2$. As a result, the equilibrium results from the baseline game are easily extended to this modified game, including the comparative statics identified in Proposition 2. In particular, an increase in σ^1 (which increases $\lambda^1\Delta^1 + \lambda^2\Delta^2$) leads $\widetilde{z}_1^*$, p_1^*, α_1^*, and X_1^* to increase, whereas it has the opposite effects on $\widetilde{z}_2^*$, p_2^*, α_2^*, and X_2^*. By contrast, if σ^2 increases, all of these effects are reversed. The proposition below clarifies how browser-level market shares are affected by defaults.

Proposition 8. *If $\eta \leq \frac{2}{3}t$, then $\frac{\partial X_i^{j*}}{\partial \sigma^i} < 0 < \frac{\partial X_i^*}{\partial \sigma^i} < \frac{\partial X_i^{i*}}{\partial \sigma^i}$ for each i and $j \neq i$.*[27]

Here X_i^{j*} denotes the equilibrium value of X_i^j. Thus increasing σ^i leads i's market share on browser i (X_i^{i*}) to increase by more than its overall market share (X_i^*). By contrast, i's market share on browser j (X_i^{j*}) *decreases*, although the latter effect requires that η is not too large. Finally, note that the inequality in Proposition 8 implies $\frac{\partial X_i^{j*}}{\partial \sigma^j} < \frac{\partial X_i^*}{\partial \sigma^j} < 0 < \frac{\partial X_i^{i*}}{\partial \sigma^j}$. In what follows, we restrict focus to pairs (σ^1, σ^2) such that: (a) each platform is active on both browsers in equilibrium ($X_i^{i*}, X_i^{j*} > 0$ for each i) and (b) platform 1 remains the overall market leader ($X_1^* > X_2^*$), although it does not necessarily have a larger market share on both browsers (that is, we could have $X_1^{2*} < X_2^{2*}$).

[26] That is, $\widetilde{x}^k$ is the location that solves $u_1^k(\widetilde{x}^k) = u_2^k(\widetilde{x}^k)$.

[27] If $\eta > \frac{2}{3}t$, then $\frac{\partial X_i^{j*}}{\partial \sigma^i}$ could become positive.

Aggregate consumer welfare is now given by $\mathcal{W}_c \equiv \lambda^1 \mathcal{W}_c^1 + \lambda^2 \mathcal{W}_c^2$, where

$$\mathcal{W}_c^k \equiv \int_0^{X_1^k} u_1^k(x)dx + \int_{X_1^k}^1 u_2^k(x)dx. \tag{23}$$

The overall impact of platform i's default contract can be broken up into two parts: a direct effect on browser i (whose users are directly impacted by the default) and a spillover effect on browser $j \neq i$. Writing welfare as a function of (σ^1, σ^2), the direct effect of platform 1's deal is $\mathcal{W}_c^1(\sigma^1, \sigma^2) - \mathcal{W}_c^1(0, \sigma^2)$, while it is $\mathcal{W}_c^2(\sigma^1, \sigma^2) - \mathcal{W}_c^2(\sigma^1, 0)$ for platform 2's agreement. Similarly, the spillover effect generated by platform 1's agreement is $\mathcal{W}_c^2(\sigma^1, \sigma^2) - \mathcal{W}_c^2(0, \sigma^2)$, while the spillover generated by platform 2's agreement is $\mathcal{W}_c^1(\sigma^1, \sigma^2) - \mathcal{W}_c^1(\sigma^1, 0)$.

Proposition 9. *Suppose that $\eta \leq \frac{2}{3}t$.*

(i) For each i, the direct effect of i's default contract on browser i is always negative.

(ii) Platform 2's default contract always generates a positive spillover on browser 1.

(iii) Platform 1's default contract generates a negative spillover on browser 2 whenever $\sigma^1 + \sigma^2$ is sufficiently small. If $\sigma^1 + \sigma^2$ is large, the spillover can be positive.

To understand these results, it is helpful to consider how a given default contract affects the four different groups of consumers (one for each browser-search platform combination). Table 1 summarizes how the dominant platform's contract affects the utility of each group.[28] The table also specifies the masses of the different groups. Table 2 provides the same information about the laggard's default.

	Search Platform 1	**Search Platform 2**
Browser 1	effect: $\frac{\partial u_1^1}{\partial \sigma^1} < 0$ mass: $\lambda^1 X_1^{1*}$	effect: $\frac{\partial u_2^1}{\partial \sigma^1} < 0$ mass: $\lambda^1 X_2^{1*}$
Browser 2	effect: $\frac{\partial u_1^2}{\partial \sigma^1} < 0$ mass: $\lambda^2 X_1^{2*}$	effect: $\frac{\partial u_2^2}{\partial \sigma^1} > 0$ mass: $\lambda^2 X_2^{2*}$

Table 1: Effects of platform 1's default contract on different consumer groups

[28] All the details needed to fill out these tables are given in the proof of Proposition 9.

	Search Platform 1	**Search Platform 2**
Browser 1	effect: $\frac{\partial u_1^1}{\partial \sigma^2} > 0$ mass: $\lambda^1 X_1^{1*}$	effect: $\frac{\partial u_2^1}{\partial \sigma^2} < 0$ mass: $\lambda^1 X_2^{1*}$
Browser 2	effect: $\frac{\partial u_1^2}{\partial \sigma^2} < 0$ mass: $\lambda^2 X_1^{2*}$	effect: $\frac{\partial u_2^2}{\partial \sigma^2} < 0$ mass: $\lambda^2 X_2^{2*}$

Table 2: Effects of platform 2's default contract on different consumer groups

By inspection of both tables, it is clear that platform i's default agreement harms all consumers on browser i, regardless of what search platform they use. This follows from the same arguments underlying Proposition 4 in the baseline game. As for spillover effects, notice that platform i's contract always benefits consumers who use platform j on browser j ($j \neq i$). This positive spillover is key to most of the results in this section.

To make things more concrete, suppose that Google (platform 1) acquires default status on Safari (browser 1), while Bing (platform 2) is the default on Firefox (browser 2). Consider Bing's deal first. It creates switching costs on Firefox, not on Safari. But it still affects Safari users indirectly in two ways. First, the contract leads Google to act as if its quality V_1 has fallen, inducing it to behave more competitively (α_1^* falls). Google users on Safari do not experience that quality reduction (they do not incur the switching cost), but they still benefit from the reduced ad intensity level. Google users also experience a reduction in network benefits (ηX_1^* falls). But so long as η is not too large, the latter effect is outweighed by the former, leaving Google users on Safari better off overall. Symmetric arguments imply that Bing users on Safari are left worse off.[29]

The masses of the four different consumer groups are also relevant. Continuing to focus on Bing's contract, consider the masses in the top row of the tables. Proposition 8 implies that $\partial X_1^{1*}/\partial \sigma^i > 0$ for both $i = 1, 2$. This ensures that $X_1^{1*} > X_1^* > \frac{1}{2}$, implying that $\lambda^1 X_1^{1*}$ is much larger than $\lambda^1 X_2^{1*}$. In other words, Bing's contract benefits far more Safari users than it harms. This explains why the laggard's contract always generates net-positive spillovers.

However, turning to Google's contract (Table 1), the operative masses (bottom row) are much less conducive to a welfare improvement. If σ^1 and σ^2 are small, then each platform's browser-level shares are close to its aggregate share ($|X_i^{j*} - X_i^*|$ is small for all i and j). In that case, $\lambda^2 X_1^{2*}$ must be larger than $\lambda^2 X_2^{2*}$, implying that Google's contract harms more Firefox users than it benefits. This explains why the dominant platform's spillover can be positive only if σ^1 and/or σ^2 are sufficiently large.

The fact that spillover effects can be positive raises the question of whether a default contract could have a positive effect on aggregate consumer welfare. The final result of this section addresses that question.

[29] These consumers experience a negative competition effect ($h\alpha_2^*$ increases) and a positive network benefit (ηX_2^* increases). The former outweighs the latter when $\eta \leq \frac{2}{3}t$, resulting in net reduction in utility.

Proposition 10. *Suppose that $\eta \leq \frac{2}{3}t$.*

(i) Platform 1's contract always reduces aggregate consumer welfare.

(ii) If λ^2 is sufficiently small, platform 2's contract may increase aggregate consumer welfare. This becomes more likely when σ^1 is larger, when η is smaller, or when μ is more concave.

Thus, even if network effects are not large, it is possible for a default contract to benefit consumers overall. But there are two important caveats. First, only the laggard's contract can potentially raise consumer welfare—the dominant platform's contract always harms consumers overall. Second, the laggard's contract can only raise welfare if it is sufficiently "narrow" in the sense that it applies to a browser with a relatively small market share. This is captured by λ^2 being small.

Intuitively, when λ^2 is small, the positive spillover generated by the laggard's agreement affects a large majority of the consumer base. For example, in the Google-Bing example, the consumer group that benefits from Bing's deal is composed of Google users on Safari. This is by far the biggest consumer group, since Google is much more popular than Bing and Safari is much more popular than Firefox. As a result, the benefits imposed on this group may outweigh the harms imposed on the other three. By contrast, when λ^2 is large, we already know from the baseline game (which is equivalent to $\lambda^2 = 1$) that the laggard's contract will reduce welfare.

The proof gives an numerical example where the laggard's agreement raises welfare. When demand is linear (as in Example 1), this can happen only if σ^1 is positive. A positive value of σ^1 makes it easier for the laggard's deal to raise welfare, because it increases the mass of consumers who benefit from the laggard's contract. However, if μ is concave, the laggard's deal could raise aggregate consumer welfare even if $\sigma^1 = 0$. Finally, a smaller value of η makes the competition-intensifying effect of the laggard's deal (which is responsible for the positive spillover) more important relative to its impact on network effects (which detracts from the positive spillover). This too makes it easier for the laggards deal to raise consumer welfare.

4 Dynamic Effects: Laggard Entry and Investment

The foregoing analysis focuses on static competitive effects. But defaults may also have dynamic effects on entry and investment, which are important to determine how defaults affect welfare over the long run. To that end, this section introduces a simple model to study how defaults affect entry and investment by potential entrants. A later section then considers the plausibility of some additional possible dynamic effects that have been widely discussed in connection with the *Google* case (Section 5.4).

The laggard is now replaced by a potential entrant who can endogenously invest in a risky R&D project aimed at creating a new search platform. If successful, it enters the market and assumes the role of platform 2. This is captured by the following three-stage game.

Stage 1: The potential entrant's type, denoted v, is drawn from a distribution G on $[\underline{v}, \overline{v}]$, where $0 \leq \underline{v} < \overline{v}$.

Stage 2: The potential entrant chooses an R&D investment $y \geq 0$. The R&D project is successful with probability $\rho(y)$.

Stage 3: If the potential entrant was successful in stage 2, it enters and assumes the role of platform 2 with quality level $V_2 = v$. The two platforms compete according to the baseline game. If the entrant failed at stage 2, platform 1 is a monopolist.

The entrant's type, v, specifies how promising it is as a potential competitor. If its type is too low, then it could not earn a positive profit in competition, in which case there is no benefit to investing. As in the baseline game, we allow for the possibility that either platform (or neither) is the default in stage 3 competition. A few assumptions will ensure the game is well behaved. First, we assume that both firms will be active in equilibrium when $V_2 = \overline{v}$, regardless of which firm (if any) is the default. Second, the probability function $\rho : [0, \infty) \to [0, 1)$ is assumed to be twice continuously differentiable with $\rho(0) = 0$, $\rho'(y) > 0$, and $\rho''(y) < 0$. It also satisfies the Inada condition $\lim_{y \to 0} \rho'(y) = \infty$.[30] Finally, we impose $\eta \leq \frac{2}{3}t$, which ensures that entry is better for consumers than monopoly.

Proposition 11. *The entry game has a unique subgame perfect Nash equilibrium, which is characterized by a function $y^*(v)$ and a threshold $\widetilde{v} \in [\underline{v}, \overline{v})$. A potential entrant's investment is positive ($y^*(v) > 0$) if and only if $v > \widetilde{v}$. Additionally, $\partial y^*(v)/\partial v > 0$ whenever $v > \widetilde{v}$. The equilibrium probability of entry is $(1 - G(\widetilde{v}))\mathbb{E}[\rho(y^*(v))|v > \widetilde{v}]$. Defaults affect the equilibrium in the following ways:*

(i) When platform 1 is the default, $\widetilde{v}$ strictly increases and $y^*(v)$ strictly falls for all types $v > \widetilde{v}$. The probability of entry strictly decreases, as does expected consumer welfare.

(ii) When platform 2 is the default, $\widetilde{v}$ strictly decreases and $y^*(v)$ strictly increases for all types $v > \widetilde{v}$. The probability of entry strictly increases. Expected consumer welfare increases if the increase in the probability of entry is sufficiently large.

A default by the dominant platform generates two harmful dynamic effects. First, it leads to some entry deterrence. The increase in $\widetilde{v}$ means that fewer potential entrants will actually attempt to enter. Second, although sufficiently strong potential entrants will still attempt to enter, they now invest less. The problem is that investment is not as worthwhile for a potential entrant, because its expected profits (conditional on entry) have been distorted downward by the dominant platform's default.

By contrast, all of these results reverse direction when platform 2 is the default, resulting in larger investments by a larger set of types, and thus a higher likelihood of entry. Hence, allowing defaults for small search engines helps to stimulate entry and investment.

To understand the consumer welfare results, recall that in the baseline game with $\eta \leq \frac{2}{3}t$, a default agreement (by either platform) necessarily reduces *static* consumer welfare (which takes

[30] This assumption can be relaxed. When the limit is finite, this creates an additional reason why some potential entrants might optimally invest zero.

entry as given). In the entry game, this reduction in post-entry static welfare is accompanied by dynamic effects that shape the probability of entry. When the dominant platform is the default, entry and investment are suppressed, and so both the static and dynamic effects of the default are harmful to consumers. By contrast, when the entrant is the default, the dynamic effects are welfare-enhancing (entry and investment increase). These dynamic effects will dominate, leading to a net increase in expected consumer welfare, when the increase in the probability of entry is sufficiently large.

5 Extensions

5.1 Diminishing Returns to User Data

In the baseline model, network effects contribute an amount $X_i\eta$ to the perceived quality of platform i. This assumes that network benefits increase linearly with scale. But, in practice, the algorithmic learning benefits generated from user data are likely subject to diminishing returns. As Tucker (2019) notes, "most studies suggest there are, at best, concave returns to data—that is, initially data can indeed provide performance advantages, but these performance advantages quickly decline as the firm obtains more data."

To explore this possibility, we now extend the baseline model as follows. There is a threshold level of scale $\hat{X} \in (0,1)$ at which marginal network benefits fall from $\overline{\eta}$ to $\underline{\eta}$, where $\overline{\eta} \geq \underline{\eta}$. In other words, the network benefits now take the form

$$\text{Network benefits } = \begin{cases} X_i\overline{\eta}, & \text{if } X_i \leq \hat{X} \\ \hat{X}\overline{\eta} + (X_i - \hat{X})\underline{\eta}, & \text{if } X_i > \hat{X}. \end{cases}$$

This is a piecewise linear function with a kink point at $X_i = \hat{X}$. If $\hat{X}$ is smaller than X_2^* or larger than X_1^*, then nothing interesting happens in this modified game.[31] The interesting situation is where $X_2^* < \hat{X} < X_1^*$. In this case, the marginal benefits of user data are larger for the laggard than for the dominant platform.

Although the main result given in this section holds for any $\hat{X} \in (X_1^*, X_2^*)$, it will be convenient to set $\hat{X} = \frac{1}{2}$. It will also be convenient to define $\overline{\eta}$ and $\underline{\eta}$ as follows:

$$\overline{\eta} \equiv (1+r)\eta, \qquad \underline{\eta} \equiv (1-r)\eta. \tag{24}$$

Here $\eta = (\overline{\eta} + \underline{\eta})/2$ is the average of $\underline{\eta}$ and $\overline{\eta}$, which we hold fixed. The parameter $r \in [0,1]$ governs the relative magnitudes of $\overline{\eta}$ and $\underline{\eta}$, and thus captures the degree to which network benefits diminish with scale. Notice that $r = 0$ just results in the baseline model. This parameterization is helpful, because equilibrium strategies and demand levels depend only on

[31] If $X_1^* < \hat{X}$, then the game is identical to the baseline model with $\eta = \overline{\eta}$. If $\hat{X} < X_2^*$ then the game is nearly identical to the baseline with parameter value $\eta = \underline{\eta}$ (the only difference is that each platform's overall quality now includes an additional contribution of $\hat{X}\overline{\eta}$).

the average parameter value, η, not on the individual values of $\overline{\eta}$ and $\underline{\eta}$.

Assuming platform 1 obtains default status, utilities for platform 1 and 2 are, respectively,

$$\begin{aligned} u_1(x) &= V_1 + \tfrac{1}{2}\overline{\eta} + (X_1 - \tfrac{1}{2})\underline{\eta} - h\alpha_1 - tx \\ &= V_1 + r\eta + X_1\underline{\eta} - h\alpha_1 - tx \end{aligned} \tag{25}$$

$$u_2(x) = V_2 + X_2\overline{\eta} - h\alpha_2 - t(1-x) - \sigma. \tag{26}$$

Following the approach from Section 2.2, consumer-side demand for platform 1 is

$$\begin{aligned} X_1 &= \frac{\Delta + t - (\overline{\eta} - r\eta) - h(\alpha_1 - \alpha_2)}{2t - \overline{\eta} - \underline{\eta}} \\ &= \frac{1}{2} + \frac{\Delta - h(\alpha_1 - \alpha_2)}{2(t - \eta)} \end{aligned} \tag{27}$$

This is precisely the same demand function from the baseline game. It follows that the first order conditions and equilibrium strategies will be the same as in the baseline game. But utilities, and thus consumer welfare effects, are now different. Our main result is that, when there are diminishing returns to data, default agreements are more detrimental to consumers.

Proposition 12. *Assume that platform 1 is the default and let $\mathcal{W}_c(r)$ denote equilibrium consumer welfare as a function of r. If $r' > r$, then*

$$\frac{\partial \mathcal{W}_c(r')}{\partial \sigma} < \frac{\partial \mathcal{W}_c(r)}{\partial \sigma}. \tag{28}$$

This result makes intuitive sense. When there are diminishing returns to data, a small amount of user data matters much more to the laggard than to the dominant platform. Thus, when a default agreement leads some consumers to switch from platform 2 to platform 1, this may generate only a small or negligible increase in the dominant platform's quality, but it may substantially diminish the laggard's quality. This asymmetry magnifies the adverse welfare effects of default agreements. These results also make it less likely that consumers might prefer monopoly to competition, as that possibility requires that increases in market share continue to generate large quality improvements even when the platform is already very large.

5.2 Asymmetric Demand Effects and Captive Consumers

Suppose that initially there are no defaults, because consumers are presented with a choice screen. Relative to this benchmark, let $\mathcal{S}_i$ denote the measure of consumers who would switch platforms (from j to i) if platform i became the default.[32] Do defaults by the two platforms have symmetric demand effects in the sense that $\mathcal{S}_1 = \mathcal{S}_2$? As explained below, the answer is likely no, and a failure to account for this asymmetry can lead to mistaken inferences about the potential anticompetitive effects of defaults.

[32] Treating X_i^* as a function of Δ, $\mathcal{S}_1 \equiv X_1^*(V_1 - V_2 + \sigma) - X_1^*(V_1 - V_2)$ and $\mathcal{S}_2 \equiv X_2^*(V_1 - V_2 - \sigma) - X_2^*(V_1 - V_2)$.

In the baseline model, the relationship between $\mathcal{S}_1$ and $\mathcal{S}_2$ is ambiguous in general. But empirical evidence indicates that the demand response is larger when the default platform faces a relatively strong competitor. That is, $\mathcal{S}_i$ is an increasing function of $V_j - V_i$. This implies that the laggard's defaults generate larger demand effects, i.e. $\mathcal{S}_2 > \mathcal{S}_1$. This is supported by empirical evidence on the effects of choice screen mandates (which eliminated Google's defaults) across different foreign jurisdictions in which Google faces varying degrees of competition.[33]

In the baseline model, the relative magnitudes of $\mathcal{S}_1$ and $\mathcal{S}_2$ depends on details about the curvature of demand.

Proposition 13. *In the baseline model:*

(i) $\mathcal{S}_1 = \mathcal{S}_2$ *if demand is linear (i.e. if* $\mu'(\cdot)$ *is constant).*

(ii) $\mathcal{S}_2 > \mathcal{S}_1$ *if* $\mu'(\cdot)$ *is strictly log-concave.*

This can be understood as follows. Recall from Section 2.2 that a default shifts the consumer demand function $X_i = X_i(\alpha_1, \alpha_2)$ vertically upward for the default platform, while the other platform's demand shifts downward. Because aggregate consumer demand is constant ($X_1 + X_2 = 1$), these countervailing shifts must be equal in magnitude. However, in principle, this shift magnitude could depend on which platform acquires default status—for example, it could be larger when the laggard is the default. But in the baseline model this doesn't happen—the shift magnitude is always $\sigma/2(t-\eta)$, regardless of which platform becomes the default. When the model is linear, this symmetry in the shift magnitudes ensures that equilibrium demand responses are likewise symmetric, i.e. $\mathcal{S}_1 = \mathcal{S}_2$. As a result, the empirically observed asymmetry ($\mathcal{S}_2 > \mathcal{S}_1$) can arise only if demand functions are curved in the right way.

However, a simple and intuitive modification of the game makes the empirically observed asymmetry much more robust, so that it no longer requires strong assumptions about demand curvature. This is achieved by supposing that some consumers are "captives" in the sense that they never switch away from the default platform (regardless of who it is). To this end, we now add a small mass $w > 0$ of captive consumers to the model.[34] For these consumers, switching costs are always prohibitively high. As in the baseline model, there is also a unit mass of "normal" consumers who face a smaller switching cost σ. Because two types differ only in the size of their switching costs, a consumer's type doesn't matter when a choice screen is implemented, since a choice screen makes switching costs irrelevant.[35]

In this extension, the magnitudes of the demand shifts are no longer symmetric.[36] Specifically,

[33] For example, Decarolis et al. (2024) find that a choice screen generated only a tiny drop (∼1%) in Google's market share in Europe, where its strongest rival (Bing) is relatively weak. (Bing's market share prior to the choice screen was less than 1%.) But in Russia, Google faces a much stronger rival—Yandex, which had a 30% market share even before the choice screen was introduced. When the Russian choice screen was implemented, Google's share fell by ∼11%.

[34] We assume that captives are uniformly distributed along the Hotelling line, and that w is small enough to ensure both platforms are active in equilibrium.

[35] Under a choice screen, a consumer never has to switch platforms in order to pick her preferred option.

[36] See the proof of Proposition 14 for details.

a default by the laggard generates larger shifts than a default by the dominant platform.[37] As a result, the desired asymmetry in equilibrium demand responses ($\mathcal{S}_2 > \mathcal{S}_1$) now arises even when there are no strong conditions on the curvature of demand. We can verify this by showing that the asymmetry occurs even when demand is linear.

Proposition 14. *In the model with captive consumers,* $\mathcal{S}_2 > \mathcal{S}_1$ *when demand is linear.*

Failing to account for this asymmetry of demand effects can lead to erroneous inferences about the competitive significance of defaults. For example, some commentators have argued that the tiny impact of the European choice screen on Google's market share in the E.U. implies that default agreements have only trivial effects on competition. But, because demand effects are asymmetric, this small effect simply reflects the weakness of Google's European competition.

The asymmetry also leads to a subtle relationship between the static and dynamic harms created by a dominant platform's defaults: if the deals have already generated substantial dynamic harm (by suppressing rival investment), culminating in a large quality gap, the defaults will then appear to have only minor effects on static competition at the margin (in the sense that $\mathcal{S}_1$ is small). Thus, evidence that a defendant's defaults have only a small effect on current market shares may reflect that the deals have already caused significant adverse effects.

5.3 Data Sharing Remedy

[TBD]

5.4 Other Possible Dynamic Effects

In Section 4, we considered how defaults affect laggard entry and investment. But the *Google* case has provoked much speculation about other possible dynamic effects of Google's defaults. For example, perhaps Google's defaults could help to promote competition in the browser market? Or perhaps they could cause additional long-run harm to the search market by discouraging a browser developer (e.g. Apple) from creating its own search engine? This section briefly considers what the model implies about the plausibility of these possibilities.

Proposition 15. *In the baseline game:*

(i) If one platform is banned from paying for default status, this reduces the other platform's willingness to pay for default status.

(ii) Total platform profits $(\Pi_1^* + \Pi_2^*)$ *are strictly higher when platform 1 is a monopolist than when both platforms are active.*

[37] The reason is that, under a choice screen, the laggard would attract only a small portion of captives (since the dominant platform is more popular), so default status greatly increases the number of captives who use it. But the same is not true for a default by the dominant platform, since it would have attracted most captives even under a choice screen.

These results shed light on some of the dynamic questions mentioned above.

Maintaining competition in the browser market. Some search access points may rely on large default payments as a primary source of revenue. For example, Mozilla (creator of Firefox) reportedly receives the majority of its revenues from Google. Therefore, if the dominant platform were banned from paying for default status, this could potentially make it hard for some browsers or other access points to sustain profitable operations. Two results potentially support this. First, Proposition 7 implies the dominant platform is willing to pay more than the laggard. Second, from part (i) of Proposition 15, eliminating the dominant platform as a potential bidder for default status would lead other search platforms to bid less.

Discouraging browsers from integrating into search. It has been widely theorized that part of the reason Google pays Apple so much for default status on Safari is that Apple was considering launching its own search engine. Google could potentially discourage this by offering to pay Apple a sufficiently large amount to maintain its default position on Safari. Apple would presumably have much less incentive to launch its own search engine if Google would remain the default option on Safari. For this story to be plausible, it would have to be the case that Apple's integration into the search market would reduce Google's profits by more than it raises Apple's. Part (ii) suggests that this is indeed the case.[38] However, unless Apple agreed to restrict its development of a competing search engine, it is not clear that this possibility raises antitrust concerns. I return to this question in Section 6.

Search platform integration into the browser market. Google created one of the biggest search access points—the Chrome browser. Google's decision to develop Chrome may have hinged on its ability to use the browser to bolster its search business by making Google Search the default browser on Chrome. In other words, if Google were prohibited from obtaining any defaults, this would have diminished its incentive to create Chrome.

These arguments suggest that some of Google's defaults could have desirable dynamic effects. However, as discussed further in Section 6, it is not clear that any of these possible dynamic benefits would require Google to be the default on *almost all* major search access points. In other words, these arguments may caution against a categorical ban on defaults, but they do not imply that there should be no antitrust limits at all.[39]

6 Discussion

United States v. Google is perhaps the most high-stakes monopolization case since *Microsoft.* But it centers on a relatively novel type of agreement that has not yet been studied extensively in economics. This paper attempts to fill that gap. The results suggest that, under plausible

[38]Even if platforms are high differentiated (large t), competition always reduces profits in this model. This is because it is consumers who view the platforms as horizontally differentiated, but it is advertisers who account for all the platforms' revenue. Any competition on the consumer side, even if it is soft, will induce the platforms to reduce their ad intensity levels (α_i), reducing the revenues that can be extracted from advertising.

[39]By analogy, antitrust does not categorically prohibit firms from engaging in exclusive dealing. Antitrust liability kicks in only when a dominant firm's exclusive dealing becomes sufficiently broad in scope.

conditions, default agreements are likely to be anticompetitive, at least if they are sufficiently broad (in terms of the share of users directly affected by them). Defaults have a biasing effect on consumer choice, distorting competition in favor of the default platform. Network effects magnify the competitive effects of this distortion. As a result, defaults can have a significant effect on competition even if switching costs are not very large.

If network effects are sufficiently strong, consumers might prefer monopoly over competition. A dominant platform's default agreements could then benefit consumers precisely because they reduce competition. If so, Google Search might be something like a natural monopoly. However, as noted in Section 5.1, most research suggests there are diminishing returns to user data, and this makes it less likely that consumers might prefer monopoly over competition.

Assuming consumers are indeed better off under competition, a broad set of default agreements involving a dominant platform is likely to harm consumers unless it happens to prevent a counterfactual arrangement that would have been even worse. To this end, if the counterfactual involves a laggard search engine acquiring default status on the same broad set of search access points, this would likely be worse for consumers. If so, the dominant platform's agreements would technically leave consumers better off, but only in the sense that they prevent the formation of even more harmful agreements.[40]

While this paper identifies many relevant competitive effects, it cannot shed much light on the magnitudes of those effects. Thus, it cannot reveal the extent to which default agreements might have contributed to Google's dominant market share. That depends on a number of empirical questions, such as precisely how "sticky" defaults are.

Vertical restraints are treated much less harshly than horizontal ones because they are often benign or procompetitive. Consistent with this, the results indicate that default agreements (particularly by laggards) could be beneficial in some situations. A sufficiently narrow default agreement by a laggard can raise consumer welfare by forcing the dominant platform to behave more competitively. And in some cases default agreements may help to stimulate entry or investment. This fits neatly with other research suggesting that vertical restraints may play an important entry-facilitating role in network industries. For example, Lee (2013) finds that exclusive rights over video games helped to facilitate successful entry by laggard video game consoles.[41] By contrast, default agreements by the dominant platform make it harder for laggards to enter. However, as noted in Section 4, some of Google's defaults could have helped to promote entry in other markets, either by Google or by independent firms.

These findings may caution against a categorical ban on default agreements. However, this does not imply that there should be no antitrust limits at all. The plausibility of a procompetitive justification grows weaker as a defendant's default contracts become broader in scope. Indeed, none of the potential dynamic justifications for defaults would seem to require a single firm to

[40]However, if the counterfactual involves a much narrower set of default agreements by the laggard, then we would have much less reason to think that the dominant platform's contracts left consumers better off. See Section 3.

[41]The comparison to video game exclusives is apt, because much like giving default status to a laggard search engine, making a popular video game exclusive to a laggard game console could reduce static welfare by nudging consumers toward a less popular product. But because this can facilitate entry, it may actually raise welfare.

obtain default status on almost all major search access points. Thus, one simple policy option would be to treat default agreements in roughly the same way we treat exclusive dealing: lawful in moderation, but not in excess.

The rest of this section discusses additional policy issues and highlights important limitations on what this paper implies about facially similar practices in other settings.

"Nudges" and antitrust. This article does not suggest that all "nudges" are likely to raise antitrust concerns. Such biasing effects are ubiquitous and rarely raise antitrust concerns. In most cases, the biasing effect is simply unavoidable. For example, competition for groceries is biased in favor of brands that get the most prominent shelf space. This bias is not a good candidate for antitrust intervention, in part because there is nothing we can do to eliminate it—*someone* has to get the best shelf space. At best, antitrust intervention could change the direction of the bias.[42] Similar arguments apply to some modern concerns about platform "self-preferencing."[43] The bias created by search engine defaults is different. It is not unavoidable—it could be eliminated through a choice screen or similar arrangement. At minimum, the magnitude of the bias could be reduced by eliminating contractual restrictions that inflate switching costs.[44]

Default-like arrangements in other settings. Similarly, this article does not suggest that defaults or similar arrangements in other contexts are likely to be anticompetitive. Default-like restrictions on consumer choice often help to reduce manufacturing costs, and some of those savings will be passed through to consumers.[45] However, this argument does not apply in the present context. Search engines and choice screens are purely digital, so we have no reason to think that search defaults reduce the cost of producing mobile devices.

Subsidization of mobile devices? Even if search engine defaults do not lower production costs, they could still reduce device prices if Google's large payments to device makers act like a subsidy. Google has cited this possibility in defense of its agreements. However, both empirical evidence and theoretical considerations cast doubt on this. Decarolis et al. (2024) studied the effects of various choice screens implemented overseas, all of which necessarily eliminate defaults on all affected devices. They found that this change had no effect on device sales, suggesting there were no significant price effects. What's more, the way payments are structured in default agreements is not conducive to significant price effects, because most device sales will not affect the payments received by the device maker. In Appendix B.3. I discuss these and other difficulties facing the subsidization argument.

Remedy considerations. If a court finds antitrust liability, what is the right remedy? The simplest remedy would be an injunction that terminates some or all of the Google's default contracts. That would leave device makers and browsers free to enter into default agreements

[42] But it is not clear that the antitrust system is well-suited to this, as it would require authorities to ask subjective questions about which brand "deserves" the best shelf space.

[43] Many of these allegations involve a platform listing its own product first within search rankings. Here too the bias is unavoidable, as *any* ranking of items will inevitably bias consumers toward the higher-ranked items.

[44] See footnote 2.

[45] For example, Ford does not let a consumer choose what brand of tires her car comes with. She is free to replace them with her preferred brand, but that will be costly. While this creates a bias in favor of the default tire brand, it is cost-efficient for Ford to install the same tires on every car.

with smaller search engines like Bing, or to offer consumers a choice screen. Another possibility, which was discussed at trial, involves changing Google's agreements so that its payments are not contingent on default status. Google could still pay device makers a cut of the search revenue they help to generate, but the payments could not be contingent on making Google the default.

Even if the court find's Google liable, its judgment will only bind Google. Thus, broad-sweeping interventions, such as a rule requiring all device makers to offer a choice screen, would require new regulation. Such a mandate would eliminate the anticopmetitive effects of defaults, but it is not necessarily the best way to do so. As noted above, allowing some limited default contracts (particularly by laggards) could be beneficial. Another proposed regulatory intervention would involve ordering Google to share its data with smaller rivals. However, Martens (2024) finds that this may not be in consumers' interest.

Less restrictive alternatives. Even if a dominant platform's default agreement leaves consumers better off by preventing a laggard from obtaining default status, the deal could still be anticompetitive in the sense that it restrains competition by more than necessary. A procompetitive benefit does not excuse a restraint if the same benefit could have been achieved by employing a "less restrictive alternative"—a less anticompetitive (but still feasible) way of achieving the same procompetitive benefit.[46] This could potentially apply in two ways. First, if the dominant platform's default agreement includes additional provisions designed to increase switching costs, those provisions would likely be anticompetitive even if the underlying default is deemed to be acceptable. Second, in principle the dominant platform could prevent the laggard from acquiring default status without having to acquire such status for itself. For example, Google and Apple could form a parity agreement in which Apple promises not treat rival search engines more favorably than Google Search. This would prevent rivals from acquiring default status, but Apple would still be allowed to implement a choice screen or other neutral system.[47]

Large payments as an entry deterrent? As noted in Section 4, if a device maker like Apple vertically integrated into the search market, this would reduce total platform profits. This implies that Google would be willing to pay a very large amount—more than the profits Apple would earn by joining the search market—to prevent Apple from entering. If Apple agreed to limit its development of a competing search engine, this would be a horizontal restraint, and could very well raise serious antitrust concerns. However, if Apple merely agreed to make Google the default search engine, without more, then antitrust scrutiny is arguably inappropriate even if the deal happens to extinguish Apple's interest in vertically integrating. Normally, when a threat of competitive entry induces an incumbent to give more favorable terms to its trading partners, we think of this as a good thing, even if it means that no entry ends up occurring. It is not clear why we should stray from that view in the present case.

[46] See, e.g., Hemphill (2016).

[47] It is easy to extend Proposition 7 to show that platform 1 would be willing to pay more for this result than platform 2 would be willing to pay for default status. This suggests that this alternative deal would be both feasible and profitable for platform 1.

References

Simon P Anderson and Bruno Jullien. The advertising-financed business model in two-sided media markets. In *Handbook of media economics*, volume 1, pages 41–90. Elsevier, 2015.

Mark Armstrong. Competition in two-sided markets. *The RAND journal of economics*, 37(3): 668–691, 2006.

Mark Armstrong and Julian Wright. Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory*, 32:353–380, 2007.

Susan Athey and Glenn Ellison. Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3):1213–1270, 2011.

Germán Bet, Roger D Blair, and Javier D Donna. The economic rationale of united states v. google. *The Antitrust Bulletin*, 67(1):23–39, 2022.

Yongmin Chen and Marius Schwartz. Assigning default position for digital goods: Competition, regulation, and welfare. Technical report, 2024.

Jay Pil Choi, Doh-Shin Jeon, and Michael Whinston. Tying in markets with network effects. *working paper*, 2023.

Alexandre De Corniere. Search advertising. *American Economic Journal: Microeconomics*, 8 (3):156–188, 2016.

Francesco Decarolis, Muxin Li, and Filippo Paternollo. *Competition and Defaults in Online Search*. Working paper, 2024.

Michal S Gal and Daniel L Rubinfeld. The hidden costs of free goods: Implications for antitrust enforcement. *Antitrust Law Journal*, 80(3):521–562, 2016.

Daniel Garcia. Search engine competition. 2023.

Richard J Gilbert and David MG Newbery. Preemptive patenting and the persistence of monopoly. *The American Economic Review*, pages 514–526, 1982.

C Scott Hemphill. Less restrictive alternatives in antitrust law. *Colum. L. Rev.*, 116:927, 2016.

Justin P Johnson. The agency model and mfn clauses. *The Review of Economic Studies*, 84 (3):1151–1185, 2017.

Robin S Lee. Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review*, 103(7):2960–3000, 2013.

Bertin Martens. The impact of search engine data sharing on competition and consumer welfare. *European Competition Journal*, pages 1–18, 2024.

John M Newman. Antitrust in zero-price markets: Foundations. *University of Pennsylvania law review*, pages 149–206, 2015.

Volker Nocke, Martin Peitz, and Konrad Stahl. Platform ownership. *Journal of the European Economic Association*, 5(6):1130–1160, 2007.

Michael Ostrovsky. Choice screen auctions. *American Economic Review*, 113(9):2486–2505, 2023.

Martin Peitz and Markus Reisinger. The economics of internet media. In *Handbook of media economics*, volume 1, pages 445–530. Elsevier, 2015.

Michael Potuck. Report reveals android users switching to iphone at 5-year high. *9to5Mac.com*, May 2023.

Andrea Prat and Tommaso Valletti. Attention oligopoly. *American Economic Journal: Microeconomics*, 14(3):530–557, 2022.

Catherine Tucker. Online advertising and antitrust: Network effects, switching costs, and data as an essential facility. *CPI Antitrust Chronicle, April*, pages 30–35, 2019.

E Glen Weyl. A price theory of multi-sided platforms. *American Economic Review*, 100(4): 1642–1672, 2010.

E Glen Weyl and Alexander White. Let the right'one'win: Policy lessons from the new economics of platforms. *Competition Policy International*, 10(2):29–51, 2014.

Appendix A: Proofs

Proof of Proposition 1

Proof. Platform i's profits are $\Pi_i = mF(\tilde{z}_i)X_i p_i$. Viewing this as a function of $(\tilde{z}_i, p_i)$, it is straightforward to verify (using assumptions 1 and 2) that the Hessian of Π_i is negative semidefinite and that Π_i is strictly concave in each argument. We will confirm shortly that the platforms' optimal choices of $\tilde{z}_i$ and p_i are always interior (i.e. they satisfy $F(\tilde{z}_i) < 1$ and $p_i > 0$) in equilibria where both platforms are active. These facts ensure that equilibrium behavior is pinned down by first order conditions (FOCs). The FOC for $\tilde{z}_i$ is

$$f(\tilde{z}_i)X_i = \frac{h}{2(t-\eta)}F(\tilde{z}_i) \quad \Longleftrightarrow \quad X_i = \frac{h}{2(t-\eta)}\mu(\tilde{z}_i), \tag{29}$$

which is equation (12) from Proposition 1. The FOC for p_i is

$$X_i = \frac{h}{2(t-\eta)}p_i. \tag{30}$$

Combining these yields $p_i = \mu(\tilde{z}_i)$, which is equation from (12) the Proposition. Thus, any equilibrium must be characterized by (11) and (12). But we must still show existence and uniqueness. Plugging the definition of X_i (from (9)) into (30) and solving for prices, we get

$$p_1^{br} = \frac{t - \eta + \Delta - h(\tilde{z}_1 - \tilde{z}_2) + hp_2}{2h}, \quad p_2^{br} = \frac{t - \eta - \Delta + h(\tilde{z}_1 - \tilde{z}_2) + hp_1}{2h}. \tag{31}$$

We can think these as the platforms' best response functions for prices, *conditional* on the values of $\tilde{z}_1$ and $\tilde{z}_2$. These functions intersect at the following prices:

$$p_1^{\dagger} = \frac{3(t - \eta) + \Delta - h(\tilde{z}_1 - \tilde{z}_2)}{3h}, \qquad p_2^{\dagger} = \frac{3(t - \eta) - \Delta + h(\tilde{z}_1 - \tilde{z}_2)}{3h}. \tag{32}$$

Suppose that the equilibrium choices of $\tilde{z}_i$ satisfy $|\tilde{z}_1 - \tilde{z}_2| \to 0$ as $\Delta \to 0$ (we will confirm this later). Then $p_1^{\dagger}$ and $p_2^{\dagger}$ are both positive for sufficiently small Δ, and $p_1^{\dagger} = p_2^{\dagger} = \frac{t-\eta}{h}$ in the limit $\Delta \to 0$. Using (11), we must have $p_i^{\dagger} = \mu(\tilde{z}_i)$ in equilibrium. Plugging this into (32) and rearranging, we get the following pair of equations:

$$\mu(\tilde{z}_1) + \frac{\tilde{z}_1}{3} = \frac{t - \eta}{h} + \frac{\Delta}{3h} + \frac{\tilde{z}_2}{3}, \qquad \mu(\tilde{z}_2) + \frac{\tilde{z}_2}{3} = \frac{t - \eta}{h} - \frac{\Delta}{3h} + \frac{\tilde{z}_1}{3}. \tag{33}$$

These equations implicitly define best response functions for $\tilde{z}_1$ and $\tilde{z}_2$. Implicit differentiation of the first equation w.r.t. $\tilde{z}_2$ shows that the slope of platform 1's best response function is $1/(1 + 3\mu'(\tilde{z}_1)) \in (0, 1)$. A symmetric argument applies to platform 2. Since the slopes of both reaction functions lie in $(0, 1)$, there exists at most one pair $(\tilde{z}_1, \tilde{z}_2)$ satisfying both equations in (33). When $\Delta = 0$, both equations are satisfied by setting $\tilde{z}_1 = \tilde{z}_2 = \tilde{z}^0$, where $\tilde{z}^0$ is defined implicitly by $\mu(\tilde{z}^0) = \frac{t-\eta}{h}$. This is the symmetric equilibrium of the game. This equilibrium has demand levels $X_1 = X_2 = \frac{1}{2}$ and $Z_1 = Z_2 = mF(\tilde{z}^0) > 0$, so both firms are active. As Δ increases from zero, best response functions shift continuously. Thus there will still be a unique equilibrium with both firms active so long as Δ is not too large.

We can derive the upper bound Δ^{crit}. The FOCs imply that $\mu(\tilde{z}_1^*) + \mu(\tilde{z}_2^*) = \frac{2(t-\eta)}{h}$. As we will verify in the proof of Proposition 2, $\tilde{z}_1^*$ is strictly increasing in Δ and $\tilde{z}_2^*$ is strictly decreasing. (This implies that $\tilde{z}_1^* > \tilde{z}_2^*$, since they are equal when $\Delta = 0$.) Thus, both firms are active so long as $\tilde{z}_2^* > \underline{z}$. Δ^{crit} is the value of Δ at which $\tilde{z}_2^* = \underline{z}$, in which case we must have $\mu(\tilde{z}_1^*) = 2(t-\eta)/h \Leftrightarrow \tilde{z}_1^* = \tilde{z}^m \equiv \mu^{-1}\big(2(t-\eta)/h\big)$. (Assumption 3 ensures this is feasible.) The corresponding prices are $p_1^* = \mu(\tilde{z}^m) = 2(t-\eta)/h$ and $p_2^* = \mu(\underline{z}) = 0$. Plugging these strategy values into (32) we find

$$\Delta^{crit} = 3(t - \eta) + h\mu^{-1}\left(\frac{2(t - \eta)}{h}\right) - h\underline{z}. \tag{34}$$

□

Proof of Proposition 2

Proof. **Part (i).** In the proof of Proposition 1, we showed that the equilibrium prices must satisfy

$$p_1^* = \frac{3(t-\eta) + \Delta - h(\tilde{z}_1^* - \tilde{z}_2^*)}{3h}, \quad p_2^* = \frac{3(t-\eta) - \Delta + h(\tilde{z}_1^* - \tilde{z}_2^*)}{3h}. \tag{35}$$

This implies $p_1^* - p_2^* = [2\Delta - 2h(\tilde{z}_1^* - \tilde{z}_2^*)]/3h$. Plugging this into the definition of X_i and simplifying, we find that equilibrium consumer demand levels must satisfy

$$X_1^* = \frac{1}{2} + \frac{\Delta - h(\tilde{z}_1^* - \tilde{z}_2^*)}{6(t-\eta)}, \qquad X_2^* = \frac{1}{2} - \frac{\Delta - h(\tilde{z}_1^* - \tilde{z}_2^*)}{6(t-\eta)}. \tag{36}$$

Substituting this into (12) for each $i = 1, 2$, we find

$$\frac{1}{2} + \frac{\Delta - h(\tilde{z}_1^* - \tilde{z}_2^*)}{6(t-\eta)} = \frac{h}{2(t-\eta)}\mu(\tilde{z}_1^*), \qquad \frac{1}{2} - \frac{\Delta - h(\tilde{z}_1^* - \tilde{z}_2^*)}{6(t-\eta)} = \frac{h}{2(t-\eta)}\mu(\tilde{z}_2^*). \tag{37}$$

Let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$, and let ∂_Δ be the derivative operator $\partial_\Delta = \frac{\partial}{\partial \Delta}$. Differentiating the equations in (37) w.r.t. Δ and rearranging, we find

$$\partial_\Delta \tilde{z}_1^* = \frac{1 + h\partial_\Delta \tilde{z}_2^*}{h(1 + 3\mu_1')}, \qquad \partial_\Delta \tilde{z}_2^* = \frac{-1 + h\partial_\Delta \tilde{z}_1^*}{h(1 + 3\mu_2')}. \tag{38}$$

These equations are linear in $\partial_\Delta \tilde{z}_1^*$ and $\partial_\Delta \tilde{z}_2^*$ and they are linearly independent, so there is a unique solution $(\partial_\Delta \tilde{z}_1^*, \partial_\Delta \tilde{z}_2^*)$. Substituting one equation into the other yields the solution

$$\partial_\Delta \tilde{z}_1^* = \frac{\mu_2'}{hM} > 0, \qquad \partial_\Delta \tilde{z}_2^* = -\frac{\mu_1'}{hM} < 0, \tag{39}$$

where $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. By differentiating (11) and (12) w.r.t. Δ, we obtain

$$\partial_\Delta p_1^* = \mu_1' \partial_\Delta \tilde{z}_1^* = \frac{\mu_1'\mu_2'}{hM} > 0, \qquad \partial_\Delta p_2^* = \mu_2' \partial_\Delta \tilde{z}_2^* = -\frac{\mu_1'\mu_2'}{hM} < 0 \tag{40}$$

and

$$\partial_\Delta X_1^* = \frac{h}{2(t-\eta)}\mu_1' \partial_\Delta \tilde{z}_1^* = \frac{\mu_1'\mu_2'}{2(t-\eta)M} > 0, \qquad \partial_\Delta X_2^* = \frac{h}{2(t-\eta)}\mu_2' \partial_\Delta \tilde{z}_2^* = -\frac{\mu_1'\mu_2'}{2(t-\eta)M} < 0. \tag{41}$$

Using the fact that $\alpha_i = \tilde{z}_i + p_i$, we have

$$\partial_\Delta \alpha_1^* = \frac{\mu_2' + \mu_1'\mu_2'}{hM} > 0, \qquad \partial_\Delta \alpha_2^* = -\frac{\mu_1' + \mu_1'\mu_2'}{hM} < 0. \tag{42}$$

Finally, using the above results, it is easy to verify that

$$\partial_\Delta [X_1^*\alpha_1^* + X_2^*\alpha_2^*] = \frac{\mu_1'\mu_2'}{hM}\left\{\frac{h}{2(t-\eta)}\left(\alpha_1^* - \alpha_2^* + \frac{\mu(\tilde{z}_1^*)}{\mu_1'} - \frac{\mu(\tilde{z}_2^*)}{\mu_2'}\right) + X_1^* - X_2^*\right\} > 0$$

where the inequality follows from Assumption 2 and the facts that $X_1^* > X_2^*$ and $\alpha_1^* > \alpha_2^*$.

Part (ii). For $\tilde{z}_i^*$, p_i^*, and α_i^* we can follow the same approach from part (i). Let ∂_η be the derivative operator $\partial_\eta = \frac{\partial}{\partial \eta}$. As in the last proof, we differentiate the equations in (37), but this time w.r.t. η. As before, this leaves us with two independent equations that are linear in the derivatives we wish to calculate (which are now $\partial_\eta \tilde{z}_i^*$ for $i = 1, 2$). Solving this system of equations gives us

$$\partial_\eta \tilde{z}_1^* = -\frac{3\mu_2' + 2}{hM} < 0, \qquad \partial_\eta \tilde{z}_2^* = -\frac{3\mu_1' + 2}{hM} < 0. \tag{43}$$

Differentiating (11) w.r.t. η, we find

$$\partial_\eta p_1^* = \mu_1' \partial_\eta \tilde{z}_1^* = -\frac{3\mu_1'\mu_2' + 2\mu_1'}{hM} < 0, \qquad \partial_\eta p_2^* = \mu_2' \partial_\eta \tilde{z}_2^* = -\frac{3\mu_1'\mu_2' + 2\mu_2'}{hM} < 0. \tag{44}$$

It follows from these results that $\partial_\eta \alpha_i^* < 0$ for each $i = 1, 2$. To finish the proof, we need to show that $\partial_\eta X_1^* = -\partial_\eta X_2^* > 0$. To do this, we show that the ratio X_1^*/X_2^* is strictly increasing in η. To this end, plugging $\mu(\tilde{z}_1^*) + \mu(\tilde{z}_2^*) = \frac{2(t-\eta)}{h}$ into (12), we get the following expression for consumer-side market shares:

$$X_i^* = \frac{\mu(\tilde{z}_i^*)}{\mu(\tilde{z}_i^*) + \mu(\tilde{z}_j^*)} \tag{45}$$

where $j \neq i$. Differentiating this identity for $i = 1$ and using results from the proof of Proposition 2, we obtain the following derivative after simplification:

$$\partial_\eta X_1^* = \frac{\mu_1'\mu_2'}{(\mu(\tilde{z}_1^*) + \mu(\tilde{z}_2^*))^2 hM} \left\{ 3(\mu(\tilde{z}_1^*) - \mu(\tilde{z}_1^*)) + 2\left(\frac{\mu(\tilde{z}_1^*)}{\mu_1'} - \frac{\mu(\tilde{z}_2^*)}{\mu_2'} \right) \right\}. \tag{46}$$

Under Assumptions 1 and 2, this is positive, which completes the proof. □

Proof of Proposition 3

Proof. Let i be the default platform and let $j = 3 - i$. As in earlier proofs, let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Using results from the proof of Proposition 2, we have

$$\begin{aligned} \frac{\partial u_j^*(x)}{\partial \sigma} &= -1 + \eta \frac{\partial X_j^*}{\partial \sigma} - h \frac{\partial \alpha_j^*}{\partial \sigma} \\ &= -1 - \frac{\eta \mu_i' \mu_j'}{2(t-\eta)M} + \frac{\mu_i' + \mu_i'\mu_j'}{M} \\ &= -\left\{ \frac{\mu_j' + 2\mu_i'\mu_j'}{M} + \frac{\eta \mu_i'\mu_j'}{2(t-\eta)M} \right\} \end{aligned}$$

which is negative, since $\mu_i', \mu_j' > 0$. Hence, users on the non-default platform are always harmed by the default. For the default platform, we have

$$\frac{\partial u_i^*(x)}{\partial \sigma} = \eta \frac{\partial X_i^*}{\partial \sigma} - h \frac{\partial \alpha_i^*}{\partial \sigma} = \frac{\eta \mu_i' \mu_j'}{2(t-\eta)M} - \frac{\mu_j' + \mu_i'\mu_j'}{M}.$$

Some algebra shows that

$$\frac{\partial u_i^*(x)}{\partial V_j} \geq 0 \quad \iff \quad \eta \leq \widehat{\eta}_i \equiv \left(\frac{2+2\mu_i'}{2+3\mu_i'}\right)t.$$

Notice that $\frac{2}{3}t < \widehat{\eta}_i < t$. Thus, if $\eta \leq \frac{2}{3}t$, the cumulative change in $u_i^*(x)$ must be negative for any x. But if η is sufficiently close to t, then the cumulative change in $u_i^*(x)$ will be positive for any x. □

Proof of Proposition 4

Proof. As shown in the proof of Proposition 5, the derivatives of $\mathcal{W}_c$ look like

$$\frac{\partial \mathcal{W}_c}{\partial V_i} = X_i^* \frac{\partial u_i^*}{\partial V_i} + X_j^* \frac{\partial u_j^*}{\partial V_i},$$

where $j \neq i$. It is easy to verify that $\partial u_i^*/\partial V_i > 0$ in all cases. In Proposition **??**, we established that the cross-derivatives $\partial u_j^*/\partial V_i$ are always positive when $\eta \leq \frac{2}{3}t$. Thus, the latter inequality implies that $\mathcal{W}_c$ is strictly increasing in V_i for each i. In this case, default agreements (involving either platform) will strictly reduce consumer welfare. As in earlier proofs, let $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Recall from Proposition 5 that

$$\begin{aligned} \frac{\partial \mathcal{W}_c}{\partial V_i} &= X_i^*\left(1 - \frac{\mu_j' + \mu_1'\mu_2'}{M} + \frac{\eta\mu_1'\mu_2'}{2(t-\eta)M}\right) + X_j^*\left(\frac{\mu_i' + \mu_1'\mu_2'}{M} - \frac{\eta\mu_1'\mu_2'}{2(t-\eta)M}\right) \\ &= \underbrace{\frac{\mu_i' + \mu_1'\mu_2'}{M} + X_i^*\frac{\mu_1'\mu_2'}{M}}_{\text{always positive}} + \underbrace{(X_i^* - X_j^*)\frac{\eta\mu_1'\mu_2'}{2(t-\eta)M}}_{\text{negative when } i=2} \end{aligned}$$

The last term in this sum diverges as $\eta \to t$. Thus, $\partial \mathcal{W}_c/\partial V_2$ is negative for sufficiently large η. Since making platform j the default is equivalent to reducing V_i by σ, this implies that a default agreement involving platform 1 will raise consumer welfare when η is sufficiently large. □

Proof of Proposition 5

Proof. We want to compare the derivatives $\frac{\partial \mathcal{W}_c}{\partial V_1}$ and $\frac{\partial \mathcal{W}_c}{\partial V_2}$. As in previous proofs, let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$, $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Using results from the proof of Proposition 2, we find

$$\begin{aligned}
\frac{\partial \mathcal{W}_c}{\partial V_2} &= \int_0^{X_1^*} \frac{\partial u_1^*(x)}{\partial V_2} dx + \int_0^{X_1^*} \frac{\partial u_1^*(x)}{\partial V_2} dx + \underbrace{u_1^*(X_1^*)\frac{\partial X_1^*}{\partial V_2} - u_2^*(X_1^*)\frac{\partial X_1^*}{\partial V_2}}_{=0,\text{ since } u_1^*(X_1^*)=u_2^*(X_1^*)} \\
&= X_1^* \frac{\partial u_1^*(x)}{\partial V_2} + X_2^* \frac{\partial u_2^*(x)}{\partial V_2} \\
&= X_1^* \left(\eta \frac{\partial X_1^*}{\partial V_2} - h \frac{\partial \alpha_1^*}{\partial V_2} \right) + X_2^* \left(1 - \eta \frac{\partial X_2^*}{\partial V_2} + h \frac{\partial \alpha_2^*}{\partial V_2} \right) \\
&= X_1^* \left(\frac{\mu_2' + \mu_1'\mu_2'}{M} - \frac{\eta \mu_1'\mu_2'}{2(t-\eta)M} \right) + X_2^* \left(1 - \frac{\mu_1' + \mu_1'\mu_2'}{M} + \frac{\eta \mu_1'\mu_2'}{2(t-\eta)M} \right).
\end{aligned}$$

A symmetric argument shows that

$$\frac{\partial \mathcal{W}_c}{\partial V_1} = X_1^* \left(1 - \frac{\mu_2' + \mu_1'\mu_2'}{M} + \frac{\eta \mu_1'\mu_2'}{2(t-\eta)M} \right) + X_2^* \left(\frac{\mu_1' + \mu_1'\mu_2'}{M} - \frac{\eta \mu_1'\mu_2'}{2(t-\eta)M} \right).$$

With these results, we can check the inequality $\frac{\partial \mathcal{W}_c}{\partial V_1} > \frac{\partial \mathcal{W}_c}{\partial V_2}$. After simplification, we find

$$\frac{\partial \mathcal{W}_c}{\partial V_1} \geq \frac{\partial \mathcal{W}_c}{\partial V_2} \quad \Longleftrightarrow \quad (X_1^* - X_2^*) \left[\frac{t}{t-\eta} \right] \mu_1'\mu_2' \geq \mu_2' - \mu_1'.$$

The lefthand side is positive and diverges as $\eta \to t$. Thus, the inequality is definitely true if $\mu(\cdot)$ is convex (implying the righthand side is nonpositive) or if η is sufficiently large. However, if $\mu(\cdot)$ is concave and η is small, then the inequality can fail. As an example, one can verify computationally that the inequality usually fails when $\mu(z) = \ln(1+z)$, $\eta = 0$, and $\frac{h}{2t} \leq 1$.

□

Proof of Proposition 6

Proof. A type-z advertiser's payoff from joining platform i is $\pi_i^*(z) = (\alpha_i^* - p_i^* - z)X_i^* = (\tilde{z}_i^* - z)X_i^*$. Thus, aggregate advertiser profits are

$$\begin{aligned}
\mathcal{W}_a &= \int_{\underline{z}}^{\tilde{z}_1^*} (\tilde{z}_1^* - z) X_1^* f(z) dz + \int_{\underline{z}}^{\tilde{z}_2^*} (\tilde{z}_2^* - z) X_2^* f(z) dz \\
&= \tilde{z}_1^* X_1^* F(\tilde{z}_1^*) - X_1^* \int_{\underline{z}}^{\tilde{z}_1^*} z f(z) dz + \tilde{z}_2^* X_2^* F(\tilde{z}_2^*) - X_2^* \int_{\underline{z}}^{\tilde{z}_2^*} z f(z) dz. \qquad (47)
\end{aligned}$$

Applying integration by parts to these integrals:

$$\int_{\underline{z}}^{\widetilde{z}_i^*} zf(z)dz = \widetilde{z}_i^* F(\widetilde{z}_i) - \int_{\underline{z}}^{\widetilde{z}_i^*} F(z)dz.$$

Plugging this into (47) and simplifying, we get

$$\mathcal{W}_a = X_1^* \int_{\underline{z}}^{\widetilde{z}_1^*} F(z)dz + X_2^* \int_{\underline{z}}^{\widetilde{z}_2^*} F(z)dz. \tag{48}$$

To prove this proposition, it is sufficient to show that $\frac{\partial \mathcal{W}_a}{\partial \Delta} > 0$. As in previous proofs, let $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Differentiating w.r.t. Δ, we find

$$\begin{aligned}
\frac{\partial \mathcal{W}_a}{\partial \Delta} &= \frac{\partial X_1^*}{\partial \Delta} \int_{\underline{z}}^{\widetilde{z}_1^*} F(z)dz + X_1^* F(\widetilde{z}_1^*) \frac{\partial \widetilde{z}_1^*}{\partial \Delta} + \frac{\partial X_2^*}{\partial \Delta} \int_{\underline{z}}^{\widetilde{z}_2^*} F(z)dz + X_2^* F(\widetilde{z}_2^*) \frac{\partial \widetilde{z}_2^*}{\partial \Delta} \\
&= \frac{\partial X_1^*}{\partial \Delta} \left(\int_{\underline{z}}^{\widetilde{z}_1^*} F(z)dz - \int_{\underline{z}}^{\widetilde{z}_2^*} F(z)dz \right) + X_1^* F(\widetilde{z}_1^*) \frac{\mu_2'}{hM} - X_2^* F(\widetilde{z}_2^*) \frac{\mu_1'}{hM} \\
&= \frac{\mu_1'\mu_2'}{2(t-\eta)M} \left(\int_{\widetilde{z}_2^*}^{\widetilde{z}_1^*} F(z)dz + \frac{\mu(\widetilde{z}_1^*)}{\mu_1'} F(\widetilde{z}_1^*) - \frac{\mu(\widetilde{z}_2^*)}{\mu_2'} F(\widetilde{z}_2^*) \right),
\end{aligned}$$

where the second and third lines use results from the proof of Proposition 2 and the third line also uses (12) to substitute $X_i^* = \frac{h}{2(t-\eta)} \mu(\widetilde{z}_i^*)$. Assumptions 1-2 imply the RHS is positive. □

Proof of Proposition 7

Proof. Technically, a platform's willingness to pay for default status depends on whether it believes the alternative involves its rival acquiring default status or neither platform obtaining default status. However, if total platform profits ($\sum_i \Pi_i^*$) are strictly increasing in Δ (which we will now prove), then the proposition is true either way. Using (11) and (12), platform i's equilibrium profits can be written as

$$\Pi_i^* = Z_i^* X_i^* p_i^* = \frac{mh}{2(t-\eta)} F(\widetilde{z}_i^*) \mu(\widetilde{z}_i^*)^2.$$

As in previous proofs, let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$, $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$, and $\partial_\Delta \equiv \frac{\partial}{\partial \Delta}$. Differentiating total platform profits w.r.t. Δ yields

$$\begin{aligned}
\partial_\Delta \sum_i \Pi_i^* &= \frac{mh}{2(t-\eta)} \sum_i \Big\{ f(\tilde{z}_i^*)\mu(\tilde{z}_i^*)^2 + 2F(\tilde{z}_i^*)\mu(\tilde{z}_i^*)\mu_i' \Big\} \partial_\Delta \tilde{z}_i^* \\
&= \frac{mh}{2(t-\eta)} \sum_i \Big\{ \mu(\tilde{z}_i^*) + 2\mu(\tilde{z}_i^*)\mu_i' \Big\} F(\tilde{z}_i^*) \partial_\Delta \tilde{z}_i^* \qquad \text{(using } f = \tfrac{F}{\mu}\text{)} \\
&= \frac{m}{2(t-\eta)M} \Big\{ \Big(\mu(\tilde{z}_1^*) + 2\mu(\tilde{z}_1^*)\mu_1' \Big) F(\tilde{z}_1^*)\mu_2' - \Big(\mu(\tilde{z}_2^*) + 2\mu(\tilde{z}_2^*)\mu_2' \Big) F(\tilde{z}_2^*)\mu_1' \Big\} \qquad \text{(using (39))} \\
&= \frac{m\mu_1'\mu_2'}{2(t-\eta)M} \left\{ \left(\frac{\mu(\tilde{z}_1^*)}{\mu_1'} + 2\mu(\tilde{z}_1^*) \right) F(\tilde{z}_1^*) - \left(\frac{\mu(\tilde{z}_2^*)}{\mu_2'} + 2\mu(\tilde{z}_2^*) \right) F(\tilde{z}_2^*) \right\}.
\end{aligned}$$

This is positive, because $\tilde{z}_1^* > \tilde{z}_2^*$ and $\left(\frac{\mu(z)}{\mu'(z)} + 2\mu(z) \right) F(z)$ is a strictly increasing function. (The latter is implied by Assumptions 1 and 2.) This completes the proof. □

Proof of Proposition 8

Proof. As in previous proofs, let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Using the results from the proof of Proposition 2, we have

$$\frac{\partial X_i^*}{\partial \sigma_i} = \frac{\lambda^i \mu_1'\mu_2'}{2(t-\eta)M} > 0, \qquad \frac{\partial X_i^*}{\partial \sigma_j} = -\frac{\lambda^j \mu_1'\mu_2'}{2(t-\eta)M} < 0$$

for each i and $j \neq i$. Notice that the definitions of X_i^i, and X_i^j imply

$$X_i^i = X_i + \frac{\lambda^j(\sigma^1 + \sigma^2)}{2t}, \qquad X_i^j = X_i - \frac{\lambda^i(\sigma^1 + \sigma^2)}{2t}. \tag{49}$$

These identities hold in equilibrium, too. Differentiating, we find

$$\frac{\partial X_i^{i*}}{\partial \sigma^i} = \frac{\lambda^i \mu_1'\mu_2'}{2(t-\eta)M} + \frac{\lambda^j}{2t}, \qquad \frac{\partial X_i^{j*}}{\partial \sigma^i} = \frac{\lambda^i \mu_1'\mu_2'}{2(t-\eta)M} - \frac{\lambda^i}{2t}.$$

Clearly $\frac{\partial X_i^{i*}}{\partial \sigma^i} > 0$. And a bit of algebra shows that

$$\frac{\partial X_i^{j*}}{\partial \sigma^i} < 0 \quad \Longleftrightarrow \quad \eta < \left(\frac{M - \mu_i'\mu_j'}{M} \right) t.$$

Since $\eta \leq \frac{2}{3}t$ (by assumption) and $\frac{M - \mu_i'\mu_j'}{M} > \frac{2}{3}$, this ensures $\frac{\partial X_i^{j*}}{\partial \sigma^i} < 0$. □

6.0.1 Proof of Proposition 9

Proof. As in previous proofs, let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. For each i and $j \neq i$, define

$$\psi_i \equiv \frac{\partial u_i^1}{\partial V_j} = \frac{\partial u_i^2}{\partial V_j} = \frac{\mu_j' + \mu_i'\mu_j'}{M} - \frac{\eta\mu_i'\mu_j'}{2(t-\eta)M}.$$

By inspection, $\psi_i < 1$. In the proof of Proposition 3, we showed that $\eta \leq \frac{2}{3}t$ implies $\psi_i > 0$. In the extension being considered, it is straightforward to verify that default agreements have the following marginal effects on consumer utility:

$$\frac{\partial u_i^i}{\partial \sigma^i} = -\lambda^i\psi_i \qquad \frac{\partial u_j^i}{\partial \sigma^i} = -(1-\lambda^i\psi_j) \tag{50}$$

$$\frac{\partial u_i^j}{\partial \sigma^i} = -\lambda^i\psi_i \qquad \frac{\partial u_j^j}{\partial \sigma^i} = \lambda^i\psi_j \tag{51}$$

The effect on the bottom right is positive, but the rest are negative. Part (i) follows from the fact that both effects in the first line are negative. To prove the other parts, we first show that $X_1^*\psi_1 > X_2^*\psi_2$. We have:

$$\begin{aligned} X_1^*\psi_1 - X_2^*\psi_2 &= X_1^*\frac{\mu_2' + \mu_1'\mu_2'}{M} - X_2^*\frac{\mu_1' + \mu_1'\mu_2'}{M} - (X_1^* - X_2^*)\frac{\eta\mu_1'\mu_2'}{2(t-\eta)M} \\ &= \frac{\mu_1'\mu_2'}{M}\left(\frac{X_1^*}{\mu_1'} - \frac{X_2^*}{\mu_2'} + (X_1^* - X_2^*)\left[\frac{2t-3\eta}{2(t-\eta)}\right]\right) \\ &= \frac{h\mu_1'\mu_2'}{2(t-\eta)M}\left(\frac{\mu(\tilde{z}_1^*)}{\mu_1'} - \frac{\mu(\tilde{z}_2^*)}{\mu_2'} + (\mu(\tilde{z}_1^*) - \mu(\tilde{z}_2^*))\left[\frac{2t-3\eta}{2(t-\eta)}\right]\right) \end{aligned} \tag{52}$$

where the third line follows from (12). Assumptions 1 and 2 ensure that $\mu(\tilde{z}_1^*) > \mu(\tilde{z}_2^*)$ and $\frac{\mu(\tilde{z}_1^*)}{\mu_1'} > \frac{\mu(\tilde{z}_2^*)}{\mu_2'}$. Thus, the expression in (52) is guaranteed to be positive when $\eta \leq \frac{2}{3}t$. Differentiating $\mathcal{W}_c^j$ w.r.t. σ^i ($j \neq i$), we find

$$\frac{\partial \mathcal{W}_c^j}{\partial \sigma^i} = X_i^{j*}\frac{\partial u_i^j}{\partial \sigma^i} + X_j^{j*}\frac{\partial u_j^j}{\partial \sigma^i} = \lambda^i\left[X_j^{j*}\psi_j - X_i^{j*}\psi_i\right]$$

Because $X_1^{1*} > X_1^*$ and $X_2^{1*} < X_2^*$ (using Proposition 8), this implies

$$\frac{\partial \mathcal{W}_c^1}{\partial \sigma^2} = \lambda^2\left[X_1^{1*}\psi_1 - X_2^{1*}\psi_2\right] > \lambda^2\left[X_1^*\psi_1 - X_2^*\psi_2\right] > 0,$$

where the final inequality follows from (52). This establishes part (ii). As for platform 1's spillover, it is

$$\frac{\partial \mathcal{W}_c^2}{\partial \sigma^1} = \lambda^2\left[X_2^{2*}\psi_1 - X_1^{2*}\psi_2\right] = \lambda^2\left[X_2^*\psi_2 - X_1^*\psi_1 + (\psi_1 + \psi_2)\frac{\lambda^1(\sigma^1 + \sigma^2)}{2t}\right],$$

where the second equality uses (49). Since $X_1^*\psi_1 > X_2^*\psi_2$, this expression can be positive only if $\sigma^1 + \sigma^2$ is sufficiently large. This establishes part (iii). □

Proof of Proposition 10

Proof. As in previous proofs, let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Let ψ_i be as defined in the proof of Proposition 9, and recall that $\psi_i > 0$ when $\eta \leq \frac{2}{3}t$. Using results from the previous proof, we obtain

$$\begin{aligned}
\frac{\partial \mathcal{W}_c}{\partial \sigma^i} &= \lambda^i \left[X_i^{i*} \frac{\partial u_i^i}{\partial \sigma^i} + X_j^{i*} \frac{\partial u_j^i}{\partial \sigma^i} \right] + \lambda^j \left[X_i^{j*} \frac{\partial u_i^j}{\partial \sigma^i} + X_j^{j*} \frac{\partial u_j^j}{\partial \sigma^i} \right] \\
&= \lambda^i \Big[-\lambda^i X_i^{i*} \psi_i - X_j^{i*}(1 - \lambda^i \psi_j) \Big] + \lambda^j \Big[-\lambda^i X_i^{j*} \psi_i + \lambda^i X_j^{j*} \psi_j \Big] \\
&= \lambda^i \Big[-\lambda^i X_i^{i*} \psi_i - X_j^{i*}(1 - \lambda^i \psi_j) - \lambda^j X_i^{j*} \psi_i + \lambda^j X_j^{j*} \psi_j \Big] \\
&= \lambda^i \Big[-X_j^{i*} + X_j^* \psi_j - X_i^* \psi_i \Big].
\end{aligned}$$

Because $X_1^* \psi_1 > X_2^* \psi_2$ (see the previous proof), it follows that

$$\frac{\partial \mathcal{W}_c}{\partial \sigma^1} = \lambda^1 \Big[-X_2^{1*} + X_2^* \psi_2 - X_1^* \psi_1 \Big] < 0.$$

This establishes part (i). For the laggard's contract we have

$$\frac{\partial \mathcal{W}_c}{\partial \sigma^2} = \lambda^2 \Big[-X_1^{2*} + X_1^* \psi_1 - X_2^* \psi_2 \Big]. \tag{53}$$

If $\sigma_1 = \sigma_2 = 0$, then $X_1^{2*} = X_1^*$, in which case the fact that $\psi_i < 1$ implies that this expression is negative. However, as σ^1 and/or σ^2 increase, X_1^{2*} falls, and the expression becomes positive when $X_1^{2*} < X_1^* \psi_2 - X_2^* \psi_2$. Thus, the overall welfare effect, $\mathcal{W}_c(\sigma^1, \sigma^2) - \mathcal{W}_c(\sigma^1, 0)$ could be positive if σ^1 is positive or if $\sigma^1 = 0$ and σ^2 is sufficiently large. We use a constructive proof to show that this is indeed possible. Consider the linear model from Example 1, and let $\eta = .4t$. Using the example, along with definitions in Section 3, this implies

$$X_1^* = \frac{1}{2} + \frac{\lambda^1 \Delta^1 + \lambda^2 \Delta^2}{6t} = \frac{1}{2} + \frac{\Delta^0 + \lambda^1 \sigma^1 - \lambda^2 \sigma^2}{6t},$$

where $\Delta^0 \equiv V_1 - V_2 > 0$. Using (49), the browser-level market shares are

$$X_1^{1*} = \frac{1}{2} + \frac{\Delta^0 + (1 + 2\lambda^2)\sigma^1 + 2\lambda^2 \sigma^2}{6t}, \qquad X_2^{1*} = 1 - X_1^{1*}$$

on browser 1 and

$$X_1^{2*} = \frac{1}{2} + \frac{\Delta^0 - 2\lambda^1 \sigma^1 - (1 + 2\lambda^1)\sigma^2}{6t}, \qquad X_2^{2*} = 1 - X_1^{2*}$$

on browser 2. It is also easy to verify that $\psi_1 = \psi_2 = \frac{1}{3}$. Thus, after some algebra, we find

$$\frac{\partial \mathcal{W}_c}{\partial \sigma^2} = \lambda^2 \Big[-X_1^{2*} + X_1^* \psi_1 - X_2^* \psi_2 \Big] = \frac{\lambda^2}{18t} \Big[8\lambda^1 \sigma^1 + (1 + 8\lambda^1)\sigma^2 - 9t - \Delta^0 \Big]. \tag{54}$$

Conditional on σ^1, let $\widetilde{\sigma}^2(\sigma^1)$ be the value of σ^1 that sets this derivative equal to zero. Then

$$\widetilde{\sigma}^2(\sigma^1) = \frac{9t + \Delta^0}{1 + 8\lambda^1} - \left(\frac{8\lambda^1}{1 + 8\lambda^1}\right)\sigma^1.$$

Now define $\widehat{\sigma}^1$ by $\widetilde{\sigma}^2(\widehat{\sigma}^1) = 0$. Then

$$\widehat{\sigma}^1 = \frac{9t + \Delta^0}{8\lambda^1}$$

If $\sigma^1 \geq \widehat{\sigma}^1$, then any default contract by the laggard (i.e. any $\sigma^2 > 0$) will increase aggregate consumer welfare. But we must confirm that this can happen in an interior equilibrium. Writing shares as a function of (σ^1, σ^2), this requires that $X_1^{1*}(\widehat{\sigma}^1, 0) < 1$ and $X_1^{2*}(\widehat{\sigma}^1, 0) > 0$. Using the above results, it is easy to verify that $X_1^{2*}(\widehat{\sigma}^1, 0) > 0$ is true for all parameter values. As for the other constraint, it is satisfied if λ^1 is sufficiently large (equivalently, if λ^2 is sufficiently small):

$$X_1^{1*}(\widehat{\sigma}^1, 0) < 1 \quad \Longleftrightarrow \quad \lambda^1 \geq \frac{9t + \Delta^0}{14t - 2\Delta^0}$$

For this inequality to be feasible, it must be that

$$\frac{9t + \Delta^0}{14t - 2\Delta^0} < 1 \quad \Longleftrightarrow \quad \Delta^0 < \tfrac{5}{3}t.$$

This constraint on t and Δ^0 is easily satisfied in an interior equilibrium. For example, if $t = \Delta^0$ and $\sigma^1 = \widehat{\sigma}^1$, any default agreement by the laggard will raise aggregate welfare iff $\lambda^1 > \frac{5}{6}$. Can a laggard default contract raise aggregate welfare if $\sigma^1 = 0$? Since the derivative $\partial \mathcal{W}_c / \partial \sigma^2$ is linear and negative when $\sigma^1 = \sigma^2 = 0$, this would require that $\sigma^2 > 2\widetilde{\sigma}^2(0)$. But we must check that this is compatible with interiority. Checking this, we find that $X_1^{1*}(0, 2\widetilde{\sigma}^2(0)) < 1$ iff $\lambda^1 > \frac{15t+5\Delta^0}{42t-4\Delta^0}$, which is feasible iff

$$\Delta^0 < 3t.$$

Separately, we find that $X_1^{2*}(0, 2\widetilde{\sigma}^2(0)) > 0$ if $\lambda^1 > \frac{5t+\Delta^0}{4\Delta^0-4t}$, which is feasible if

$$\Delta^0 > 3t.$$

These feasibility constraints are exact opposites. Thus, $\sigma^2 = 2\widetilde{\sigma}^2(0)$ is just barely too big to generate an interior equilibrium. However, such an equilibrium could be obtained if μ were concave rather than linear. In that case we would have $\psi_1 > \psi_2$ (also ψ_1 can be larger than $\frac{1}{2}$ if μ is sufficiently concave). This would allow the derivative in (53) to become positive at lower values of σ^2. A lower value of η also makes it easier for the laggard's contract to raise welfare, since this makes ψ_1 and ψ_2 larger. These points establish part (*ii*). □

Proof of Proposition 11

Proof. Given quality levels V_1 and V_2, let $\Pi_i^* = \Pi_i^*(V_1, V_2)$ denote platform i's equilibrium profits in the baseline game when there are no defaults. Fix a type $v \in [\underline{v}, \overline{v}]$. If it enters, its profit (ignoring its investment cost) would be $\Pi_2^*(V_1, v)$. This is positive iff $v_2 > \widetilde{v}$, where $\widetilde{v}$ is pinned down by

$$V_1 - \widetilde{v} = \Delta^{crit}. \tag{55}$$

Recall that Δ^{crit} is the threshold value of Δ such that platform 2 is inactive whenever $\Delta \geq \Delta^{crit}$. Thus, types $v \leq \widetilde{v} = V_1 - \Delta^{crit}$ will not invest anything ($y^*(v) = 0$). Now fix a type $v > \widetilde{v}$. Conditional on investing y, its ex-ante expected profits are

$$\rho(y)\Pi_2^*(V_1, v) - y.$$

Since $\rho''(y) < 0$, this is strictly concave in y. The Inada condition $\lim_{y \to 0} \rho'(y) = \infty$ ensures that the optimal investment, y^*, is a positive number that solves the first order condition

$$\rho'(y^*)\Pi_2^*(V_1, v) - 1 = 0. \tag{56}$$

Treating the solution as a function $y^* = y^*(v)$, we can totally differentiate this equation w.r.t. v and rearrange to get

$$\frac{\partial y^*(v)}{\partial v} = -\frac{\rho'(y^*(v))\frac{\partial \Pi_2^*(V_1,v)}{\partial v}}{\rho''(y^*(v))\Pi_2^*(V_1, v)} > 0. \tag{57}$$

This shows that the optimal investment $y^*(v)$ is strictly increasing in v whenever the optimal investment is positive.

When platform 1 is the default, the effects on the entrants payoffs are equivalent to simply reducing v by σ. In this case, the threshold $\widetilde{v}$ increases to $\widetilde{v}^+ \equiv V_1 + \sigma - \Delta^{crit}$. Additionally, every type $v > \widetilde{v}^+$ now invests less: it now behaves as if its type is $v - \sigma$, which induces a smaller investment, since $\partial y^*/\partial v > 0$. These effects obviously imply that the probability of entry strictly decreases. By contrast, if platform 2 is the default, then these effects are reversed, so the probability of entry increases.

Because $\eta \leq \frac{2}{3}t$, we know that defaults by either firm strictly reduce static consumer welfare. It also implies that consumers are better off with a competitive laggard than with platform 1 being a monopolist. That is, $\mathcal{W}_c(V_1, V_2) > \mathcal{W}_c^M$ for any quality levels (V_1, V_2) such that both firms are active in the baseline game. Here $\mathcal{W}_c = \mathcal{W}_c(V_1, V_2)$ is the static consumer welfare function from the main text, and $\mathcal{W}_c^M$ gives consumer welfare when platform 1 is a monopolist. It turns out that $\mathcal{W}_c^M = \frac{1}{2}t$.[48] Expected consumer welfare is

$$G(\widetilde{v})\mathcal{W}_c^M + \int_{\widetilde{v}}^{\overline{v}} \left\{\rho(y^*(v))\mathcal{W}_c(V_1, v) + [1 - \rho(y^*(v))]\mathcal{W}_c^M\right\} dG(v).$$

Under a default agreement, the post-entry static welfare level for any type v falls from $\mathcal{W}_c(V_1, v)$ to either $\mathcal{W}_c(V_1, v - \sigma)$ or $\mathcal{W}_c(V_1 - \sigma, v)$. This effect, taken alone, reduces expected consumer welfare. But there is also an effect on the probability of entry, as noted above. If platform 1 is the default, the probability falls, further reducing expected consumer welfare. Hence, expected consumer welfare definitely falls when platform 1 is the default. By contrast, when platform

[48]When platform 1 is a monopolist, it sets α_1 so that $u_1(1) = 0$, i.e. the consumer at location $x = 1$ gets exactly zero utility from joining platform 1. This is the largest α_1 that ensures all consumers join platform 1. This implies that a consumer's equilibrium utility is just $u_1(x) = (1 - x)t$. This leads to an aggregate consumer welfare level of $\frac{1}{2}t$. Hence consumer welfare does not depend on the monopolist's quality (V_1).

2 is the default, the probability of entry increases. In this case, the net change in expected consumer welfare is in general ambiguous. If the increase in the probability of entry is large enough, then it will increase.

For example, because platform 2's default reduces $\tilde{v}$ to $\tilde{v}^- \equiv V_1 - \sigma - \Delta^{crit}$, it will necessarily raise the welfare contributions from all types $v \in (\tilde{v}^-, \tilde{v} + \varepsilon)$ for sufficiently small $\varepsilon > 0$. All of these types contribute positive amounts to expected consumer welfare. But absent the default, types $v \in (\tilde{v}^-, \tilde{v}]$ would contribute nothing to welfare and types $v \in (\tilde{v}, \tilde{v} + \varepsilon)$ would contribute very little, since $y^*(v)$ would be close to zero in a right-neighborhood of $v = \tilde{v}$. Now suppose that G assigns almost all probability to types in $v \in (\tilde{v}^-, \tilde{v} + \varepsilon)$. Then the default would obviously increase expected consumer welfare. This is because the dynamic increase in the probability of entry is especially large, since the most likely types are those who make nontrivial investments only when platform 2 is the default. By contrast, if probability density is concentrated mainly on high types who are already very likely to enter *even without* a default, then the default has little effect on the probability of entry, and it would reduce expected welfare. □

Proof of Proposition 12

Proof. As we saw in the proof of Proposition 5, the derivatives of $\mathcal{W}_c$ take the form

$$\frac{\partial \mathcal{W}_c(r)}{\partial \sigma} = X_1^* \frac{\partial u_1^*}{\partial \sigma} + X_2^* \frac{\partial u_2^*}{\partial \sigma}.$$

As in previous proofs, let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Using the results from the proof of Proposition 2, we have

$$\begin{aligned}
\frac{\partial u_1^*}{\partial \sigma} &= -\left(h\frac{\partial \alpha_1^*}{\partial \Delta} - \underline{\eta}\frac{\partial X_1^*}{\partial \Delta}\right) \\
&= -\left(\frac{\mu_2' + \mu_1'\mu_2'}{M} - \frac{(1-r)\eta\mu_1'\mu_2'}{2(t-\eta)M}\right) \\
&= -\left(\frac{\mu_2' + \mu_1'\mu_2'}{M} - \frac{\eta\mu_1'\mu_2'}{2(t-\eta)M}\right) - \frac{r\eta\mu_1'\mu_2'}{2(t-\eta)M}
\end{aligned}$$

And, by a similar process, we find

$$\frac{\partial u_2^*}{\partial \sigma} = -\left(1 - \frac{\mu_1' + \mu_1'\mu_2'}{M} + \frac{\eta\mu_1'\mu_2'}{2(t-\eta)M}\right) - \frac{r\eta\mu_1'\mu_2'}{2(t-\eta)M} \tag{58}$$

Therefore

$$\frac{\partial \mathcal{W}_c(r)}{\partial \sigma} = -X_1^*\left(\frac{\mu_2' + \mu_1'\mu_2'}{M}\right) - X_2^*\left(1 - \frac{\mu_1' + \mu_1'\mu_2'}{M}\right) \tag{59}$$

$$+ (X_1^* - X_2^*)\frac{\eta\mu_1'\mu_2'}{2(t-\eta)M} - (X_1^* + X_2^*)\frac{r\eta\mu_1'\mu_2'}{2(t-\eta)M} \tag{60}$$

$$= \left.\frac{\partial \mathcal{W}_c(r)}{\partial \sigma}\right|_{r=0} - \frac{r\eta\mu_1'\mu_2'}{2(t-\eta)M} \tag{61}$$

The first term, $\frac{\partial \mathcal{W}_c(r)}{\partial \sigma}|_{r=0}$, is simply the welfare derivative from the baseline game, which corresponds to $r = 0$. Clearly the righthand side is strictly decreasing in r. □

Proof of Proposition 13

Proof. Treating X_i^* as a function of Δ, $\mathcal{S}_1 \equiv X_1^*(V_1 - V_2 + \sigma) - X_1^*(V_1 - V_2)$ and $\mathcal{S}_2 \equiv X_2^*(V_1 - V_2 - \sigma) - X_2^*(V_1 - V_2)$. It is easy to see from the linear model example (Section 2.4) that $\mathcal{S}_1 = \mathcal{S}_2$ in that case. More generally, it is clear that $\mathcal{S}_2 > \mathcal{S}_1$ is implied by X_1^* being strictly concave in Δ. Recall from the proof of proposition 2 that

$$\frac{\partial X_1^*}{\partial \Delta} = \frac{\mu_1'\mu_2'}{2(t-\eta)M},$$

where $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. It follows that X_1^* is strictly concave in Δ if and only if

$$\frac{\mu_1'\mu_2'}{\mu_1' + \mu_2' + 3\mu_1'\mu_2'} = \frac{1}{\frac{1}{\mu_2'} + \frac{1}{\mu_1'} + 3}$$

is strictly increasing, which is equivalent to the condition that $\frac{1}{\mu_1'} + \frac{1}{\mu_2'}$ is strictly increasing in Δ. Taking the derivative of this expression, and using the fact that $\frac{\partial \widetilde{z}_1^*}{\partial \Delta} = \frac{\mu_2'}{hM}$ and $\frac{\partial \widetilde{z}_2^*}{\partial \Delta} = -\frac{\mu_1'}{hM}$, we find that $\frac{1}{\mu_1'} + \frac{1}{\mu_2'}$ is strictly increasing iff

$$\phi(\widetilde{z}_1^*) < \left(\frac{\mu_1'}{\mu_2'}\right)^2 \phi(\widetilde{z}_2^*), \tag{62}$$

where $\phi(z) \equiv \frac{d\ln\mu'(z)}{dz} = \frac{\mu''(z)}{\mu'(z)}$. To complete the proof, we assume that ϕ is strictly decreasing (which is equivalent to strict log-concavity of μ') and show that this implies inequality (62). Note that $\widetilde{z}_1^* > \widetilde{z}_2^*$ and $\mu_i' > 0$ (by Assumption 1). Then the inequality is obviously true if either (a) $\phi(\widetilde{z}_1) = 0 < \phi(\widetilde{z}_2^*)$; (b) $\phi(\widetilde{z}_1) < 0 = \phi(\widetilde{z}_2^*)$; or (c) $\phi(\widetilde{z}_1^*) < 0 < \phi(\widetilde{z}_2^*)$. There are just two other possibilities. If $0 < \phi(\widetilde{z}_1^*) < \phi(\widetilde{z}_2^*)$, then it must be that μ is strictly convex ($\mu''(z) > 0$) over this range. This implies that $\frac{\mu_1'}{\mu_2'} > 1$, in which case (62) is clearly true. Finally, suppose that $\phi(\widetilde{z}_1^*) < \phi(\widetilde{z}_2^*) < 0$. This corresponds to μ being strictly concave $\mu''(z) < 0$. This implies that $\frac{\mu_1'}{\mu_2'} \in (0, 1)$. Then, keeping in mind that both $\phi(\widetilde{z}_1^*)$ and $\phi(\widetilde{z}_2^*)$ are negative in this case, this implies once again that (62) is true. □

Proof of Proposition 14

Proof. There is now a small mass $w > 0$ of captives, along with a unit mass of "normal" types. Both types are uniformly distributed. Recall that $\delta_i = 1$ if platform i is the default and $\delta_i = 0$ otherwise. We will use a superscript of $\delta_1\delta_2$ to indicate whether a default is in effect ($\delta_1\delta_2 = 10$ or $\delta_1\delta_2 = 01$) or whether a consumers face a choice screen ($\delta_1\delta_2 = 00$). Let $\widetilde{x}^{\delta_1\delta_2}$ be the marginal normal type consumer. It is pinned down by the indifference condition

$$\begin{aligned} & u_1^{\delta_1\delta_2}(\widetilde{x}^{\delta_1\delta_2}) = u_2^{\delta_1\delta_2}(\widetilde{x}^{\delta_1\delta_2}) \\ \iff \quad & V_1 - \delta_2\sigma + X_1^{\delta_1\delta_2}\eta - h\alpha_1 - t\widetilde{x}^{\delta_1\delta_2} = V_2 - \delta_1\sigma + X_2^{\delta_1\delta_2}\eta - h\alpha_2 - t(1 - \widetilde{x}^{\delta_1\delta_2}), \end{aligned} \tag{63}$$

where consumer demand levels, $X_i^{\delta_1\delta_2}$, are now given by

$$X_1^{\delta_1\delta_2} = \begin{cases} \widetilde{x}^{10} + w, & \text{if } \delta_1\delta_2 = 10 \\ \widetilde{x}^{01}, & \text{if } \delta_1\delta_2 = 01 \\ (1+w)\widetilde{x}^{00}, & \text{if } \delta_1\delta_2 = 00, \end{cases} \qquad X_2^{\delta_1\delta_2} = \begin{cases} 1 - \widetilde{x}^{10}, & \text{if } \delta_1\delta_2 = 10 \\ 1 - \widetilde{x}^{01} + w, & \text{if } \delta_1\delta_2 = 01 \\ (1+w)(1 - \widetilde{x}^{00}), & \text{if } \delta_1\delta_2 = 00. \end{cases} \tag{64}$$

This reflects the fact that, under a default, all captives subscribe to the default platform. By contrast, under a choice screen, captive and normal types are equivalent, and $\widetilde{x}^{00}$ is the marginal location for both types. Notice that $X_1^{\delta_1\delta_2} + X_2^{\delta_1\delta_2} = 1 + w$ in all cases.

Using the above expressions, when a default is in effect, the marginal location is given by one of the following two expressions:

$$\widetilde{x}^{10} = \frac{1}{2} + \frac{V_1 - V_2 + \sigma + w\eta - h(\alpha_1 - \alpha_2)}{2(t-\eta)}, \quad \widetilde{x}^{01} = \frac{1}{2} + \frac{V_1 - V_2 - \sigma - w\eta - h(\alpha_1 - \alpha_2)}{2(t-\eta)}. \tag{65}$$

By contrast, under a choice screen, the marginal location is

$$\widetilde{x}^{10} = \frac{1}{2} + \frac{V_1 - V_2 - h(\alpha_1 - \alpha_2)}{2[t - (1+w)\eta]}. \tag{66}$$

Notice that the denominator is different under a choice screen. As a result, in this extension, a default leads the demand functions to undergo both a shift and a "twist" (a change in slope). However, the twists are symmetric, and we can suppress them by imposing $\alpha_1 = \alpha_2 = 0$. This leads to the following expressions for the vertical shifts

$$\begin{aligned} X_1^{10} - X_1^{00} &= \frac{\sigma + wt}{2(t-\eta)} - \frac{(V_1 - V_2)wt}{2(t-\eta)[t - (1+w)\eta]} \\ X_2^{01} - X_2^{00} &= \frac{\sigma + wt}{2(t-\eta)} + \frac{(V_1 - V_2)wt}{2(t-\eta)[t - (1+w)\eta]}. \end{aligned}$$

The first (resp. second) expression gives the magnitude of the demand shifts when the dominant platform (resp. laggard) is the default. So long as $(1+w)\eta < t$ (which is necessary for both firms to be active), the laggard's default generates strictly larger vertical shifts.

Analogous to the baseline game, platform profits are $\Pi_i = mF(\widetilde{z}_i)X_i^{\delta_1\delta_2}p_i$. It is easy to verify that the FOC $p_i^* = \mu_i(\tilde{z}_i^*)$ continues to apply in all cases in this extension. Additionally, when a default is in effect, the second FOC is also the same: $X_i^{\delta_1\delta_2*} = \frac{h}{2(t-\eta)}\mu(\tilde{z}_i^*)$. However, under a choice screen, the second FOC changes slightly: it becomes $\tilde{x}^{00*} = \frac{h}{2(t-\eta)}\mu(\tilde{z}_1^*)$ for platform 1 and $1 - \tilde{x}^{00*} = \frac{h}{2(t-\eta)}\mu(\tilde{z}_2^*)$ for platform 2.

When demand is linear, z is drawn from the uniform distribution on $[0, \ell]$, so that $F(x) = z/\ell$ and $\mu(z) = z$. The FOC $p_i^* = \mu(\tilde{z}_i^*)$ thus implies $p_i^* = \tilde{z}_i^*$ and thus $\alpha_i^* = p_i^* + \tilde{z}_i^* = 2\tilde{z}_i^*$. Plugging these expressions into the FOCs and solving the equations, we find the following equilibrium results. First, when platform 1 is the default:

$$\tilde{x}^{10*} = \frac{5-4w}{10} + \frac{V_1 - V_2 + \sigma + w\eta}{10(t-\eta)}, \quad X_1^{10*} = \tilde{x}^{10*} + w, \quad X_2^{10*} = 1 - \tilde{x}^{10*}$$
$$\tilde{z}_1^{10*} = p_1^{10*} = \left(\frac{5+6w}{5}\right)\frac{t-\eta}{h} + \frac{V_1 - V_2 + \sigma + w\eta}{5h}$$
$$\tilde{z}_2^{10*} = p_2^{10*} = \left(\frac{5+4w}{5}\right)\frac{t-\eta}{h} - \frac{V_1 - V_2 + \sigma + w\eta}{5h}$$

When platform 2 is the default:

$$\tilde{x}^{01*} = \frac{5+4w}{10} + \frac{V_1 - V_2 - \sigma - w\eta}{10(t-\eta)}, \quad X_1^{01*} = \tilde{x}^{01*}, \quad X_2^{01*} = 1 - \tilde{x}^{01*} + w$$
$$\tilde{z}_1^{01*} = p_1^{01*} = \left(\frac{5+4w}{5}\right)\frac{t-\eta}{h} + \frac{V_1 - V_2 - \sigma - w\eta}{5h}$$
$$\tilde{z}_2^{01*} = p_2^{01*} = \left(\frac{5+6w}{5}\right)\frac{t-\eta}{h} - \frac{V_1 - V_2 - \sigma - w\eta}{5h}$$

Finally, under a choice screen:

$$\tilde{x}^{00*} = \frac{1}{2} + \frac{V_1 - V_2}{10[t-(1+w)\eta]}, \quad X_1^{00*} = (1+w)\tilde{x}^{00*}, \quad X_2^{00*} = (1+w)(1-\tilde{x}^{00*})$$
$$\tilde{z}_1^{00*} = p_1^{00*} = \frac{t-(1+w)\eta}{h} + \frac{V_1 - V_2}{5h}, \quad \tilde{z}_2^{00*} = p_2^{00*} = \frac{t-(1+w)\eta}{h} - \frac{V_1 - V_2}{5h}$$

Using these expressions, we can calculate equilibrium demand responses of defaults:

$$\mathcal{S}_1 = X_1^{10*} - X_1^{00*} = \frac{\sigma + w\eta + V_1 - V_2}{10(t-\eta)} - \frac{(1+w)(V_1 - V_2)}{10[t-(1+w)\eta]}$$
$$\mathcal{S}_2 = X_2^{10*} - X_2^{00*} = \frac{\sigma + w\eta - (V_1 - V_2)}{10(t-\eta)} + \frac{(1+w)(V_1 - V_2)}{10[t-(1+w)\eta]}$$

This implies

$$\mathcal{S}_2 - \mathcal{S}_1 = \frac{(V_1 - V_2)wt}{5(t-\eta)[t-(1+w)\eta]} > 0,$$

as desired. By inspection, this difference is proportional to the difference in the magnitudes of the vertical demand shifts (computed earlier). □

Proof of Proposition 15

Proof. **Part (i).** Let $\Pi_i^*(V_i, V_j)$ denote a platform's profits as a function of (V_i, V_j). Platform i's willingness to pay (WTP) depends on its disagreement payoff. If it believes that disagreement would lead platform j to acquire default rights, then its WTP is $\Pi_i^*(V_i, V_j - \sigma) - \Pi_i^*(V_i - \sigma, V_j)$. But if platform j is banned from acquiring default rights, then disagreement simply preserves the status quo, and hence platform i's WTP is $\Pi_i^*(V_i, V_j - \sigma) - \Pi_i^*(V_i, V_j)$. The latter WTP is strictly smaller than the former, which establishes the desired result.

Part (ii). In the proof of Proposition 7, we showed that total platform profits $(\Pi_1^* + \Pi_2^*)$ are strictly increasing in Δ. When Δ gets large enough, platform 1 becomes a monopolist. This establishes the desired result. □

Appendix B: Supplementary Material

B.1. The European Android Choice Screen

This figure shows the choice screen that Google was ordered to offer on Android devices within the European Economic Area. The order of search engines on the choice screen is randomized.

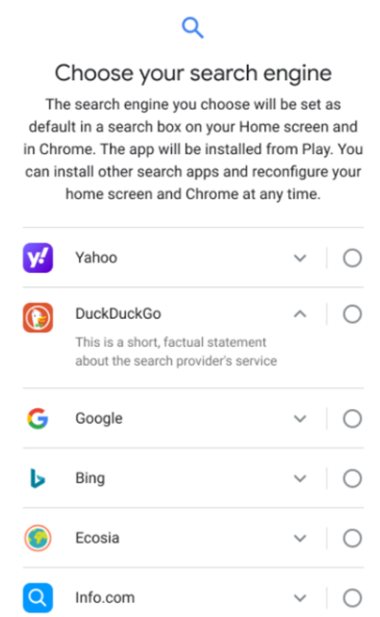

Figure 1: The Android choice screen implemented in the EEA.

B.2. Concavity/Convexity of μ

To help understand the relevance of the concavity or convexity of $\mu(\cdot)$, it is helpful to define the following locus of points:

$$\mathcal{Z}^* = \left\{ (\tilde{z}_1, \tilde{z}_2) \,\middle|\, \mu(\tilde{z}_1) + \mu(\tilde{z}_2) = \tfrac{2(t-\eta)}{h},\ \mu(\tilde{z}_1) \geq 0, \mu(\tilde{z}_2) \geq 0 \right\} \tag{67}$$

Using (11), (12), and (13), this locus contains all the different equilibria in which both platforms are active—one for each value of $\Delta \in (-\Delta^{crit}, \Delta^{crit})$.[49] It traces out a curve in the plane, and the shape of this curve is relevant to certain results in the paper. In the figure below, we plot different possible shapes of $\mathcal{Z}^*$, which corresponds to different functional forms of $\mu(z)$. Because we assume $\Delta > 0$, we are only interested in the portion of $\mathcal{Z}^*$ that lies above the 45-degree line. When $\Delta = 0$, the equilibrium is at the intersection of $\mathcal{Z}^*$ and the 45-degree line. As Δ increases, the equilibrium point travels northwest along the curve until it terminates at $(\tilde{z}^m, \underline{z})$ when $\Delta = \Delta^{crit}$. (For simplicity, the figure assumes that $\underline{z} = 0$.)

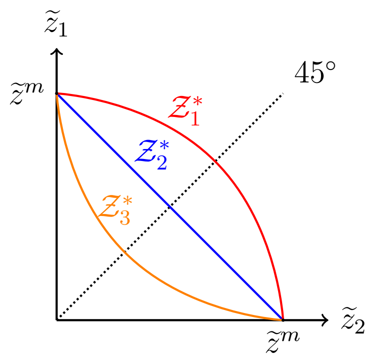

Figure 2: Examples of $\mathcal{Z}^*$ corresponding to different functional forms of $\mu(\cdot)$.

In the figure, the examples $\mathcal{Z}_1^*$, $\mathcal{Z}_2^*$, and $\mathcal{Z}_3^*$ correspond to $\mu(\cdot)$ being convex, linear, and concave, respectively. Starting at a point $\Delta \in (0, \Delta^{crit})$, a small increase in Δ always leads to an increase in $\tilde{z}_1$ and a decrease in $\tilde{z}_2$, but the relative magnitudes of these effects depend on the shape of μ. If μ is convex, then the change in $\tilde{z}_1$ is small and the change in $\tilde{z}_2$ is large, while the opposite is true when μ is concave. The same is true of the relative changes in α_1 and α_2. This is because $\alpha_i = \tilde{z}_i + p_i$, and we know from (13) that the changes in p_1 and p_2 are always equal in magnitude (i.e. $\frac{\partial p_1}{\partial \Delta} = -\frac{\partial p_2}{\partial \Delta}$).

Thus, when μ is convex, a default agreement involving the dominant firm generates a small increase in α_1 and a large reduction in α_2. But if μ is concave, then the deal generates a large increase in α_1 but only a small reduction in α_2. Since larger values of α_1 and α_2 are associated with lower consumer welfare but higher platform profits, this tells us about how the curvature of μ relates to the welfare and profit effects of default agreements. When μ is convex, a small increase in Δ is less detrimental to consumer welfare, but also less conducive to total platform profits $(\Pi_1 + \Pi_2)$. But when μ is concave, a small increase in Δ increases total platform profits significantly, but it is also especially harmful to consumers.

[49] Following (11), for a given pair $(\tilde{z}_1, \tilde{z}_2) \in \mathcal{Z}^*$, the corresponding prices are $(p_1, p_2) = (\mu(\tilde{z}_1), \mu(\tilde{z}_2))$.

B.3. Device-Level Subsidization and Pass-Through

In other industries, default-like arrangements are extremely common, and yet they usually do not arouse any controversy. One key reason for this is that such arrangements often serve to reduce production costs, and some of those savings will be passed through to consumers.[50] This could offset the adverse effects of the choice restrictions. Many of the firms Google pays for default status are mobile device makers, such as Apple, Samsung, and LG. This raises the question of whether any adverse competitive effects of search engine default agreements might be offset by price cuts at the device-level. If so, however, the mechanism driving the price cuts would have to be something other than the manufacturing economies discussed above. Search engines and choice screens are purely digital, so we have no reason to think that these default agreements reduce the cost of producing mobile devices.

An alternative possibility is that Google's large payments to device makers could act like a subsidy, inducing device makers to offer lower prices. This is a common argument for why default agreements could potentially be procompetitive. To test it directly would require a larger model encompassing both search engine competition and mobile device competition, which is beyond the scope of this paper. However, we can shed some light on it by identifying some challenges one would have to overcome to make the subsidization argument work.

First, recent empirical work finds no evidence of price effects. If default agreements lead to significant device-level price cuts, then one would expect prices to rise if a regulator proscribed default agreements and ordered device makers to offer a choice screen instead. Regulators in Europe and several other jurisdictions recently did just that for Android devices. Decarolis et al. (2024) studied the effects of this shift, and they found no evidence that it had *any* impact on device prices.

Second, the payment structure in search engine default agreements is not conducive to significant price effects. Consider the Apple agreement as an example. Google's payments to Apple are not based on iPhone sales. Rather, Google pays Apple a share of the advertising revenue it earns from user activity on Safari. For this reason, most iPhone sales have no effect on Google's payments to Apple, and hence are not "subsidized" at all. If Apple cuts the iPhone's price, it will sell more units, but most of those new sales will be upgrades by existing iPhone users.[51] Those users' search activity was *already* contributing to Apple's payments, so these iPhone sales have no effect on the default payments.

Third, the default payments give device makers an incentive to raise switching costs (e.g. by making it harder to switch browsers) to steer consumers toward Google, since this will increase their payments. This effect would exacerbate the anticompetitive effects of default agreements. Finally, note that the subsidization argument does not apply to all default agreements, because many of the firms who control search access points do not sell anything to consumers. For example, Mozilla's Firefox browser is free to consumers, so the Firefox default agreement

[50] For example, Ford does not let a consumer choose what brand of tires their car comes with. A consumer who strongly prefers a different brand will have to buy them herself, which is costly. While this creates a strong bias in favor of the default tire brand, it is cost-efficient for Ford to commit to a single brand.

[51] One recent study finds that more than 80% of iPhone sales are upgrades by existing iPhone users. Potuck (2023) summarizes the study.

obviously will not generate any offsetting price effects.

Finally, as the next result shows, even the largest possible device-level price cut may not be big enough to offset the harm to consumers. In the best-case scenario, the dominant platform would pay its full WTP for default status, and the device maker would fully pass through that amount to consumers. In such a case, whether consumers are left better off overall depends on whether the default agreement raises the dominant platform's profits by more than it reduces consumer welfare. For an infinitesimal switching cost, this boils down to the requirement that

$$\frac{\partial \Pi_1^*}{\partial \sigma} \geq -\frac{\partial \mathcal{W}_c^*}{\partial \sigma} \tag{68}$$

Assuming $\eta \leq \frac{2}{3}t$, it is straightforward to verify that the following inequality is a necessary condition for (68):

$$\frac{m}{h}[1 + 2\mu_1'(\widetilde{z}_1^*)]F(\widetilde{z}_1^*)X_1^* \geq 1 + 3\mu'(\widetilde{z}_1^*)X_1^*. \tag{69}$$

This inequatlity will fail if $\frac{m}{h}$ is not too large.