The Competitive Effects of Search Engine Defaults

Erik Hovenkamp*

Journal of Law and Economics (forthcoming, 2027)

Abstract

The landmark antitrust case *U.S. v. Google* centers on novel contracts in which Google pays browsers and device makers to make it the default search engine on their products. But the competitive effects of such arrangements are not yet well understood. I study them in a model of competition between two-sided search platforms, which earn all revenues from advertisers. Search algorithms learn and improve with use, a process that resembles network effects. Defaults create switching costs that nudge consumers toward the default search engine. 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 firm typically reduce consumer welfare. However, narrow defaults (particularly by laggards) can stimulate dynamic entry and investment while also potentially increasing static welfare. This casts doubt on proposed "choice screen" mandates that would eliminate defaults altogether.

*Cornell University. For helpful comments, I thank Ryan Bubb, Juan Castillo, Javier Donna, Josh Fischman, Rich Gilbert, Jacob Goldin, Andrei Hagiu, Tom Hubbard, Justin Johnson, Mike Katz, Doug Melamed, Eric Posner, Steve Salop, Megan Stevenson, and Mike Whinston. I'm also grateful to organizers and participants at the Northwestern University Conference on Antitrust Economics and Competition Policy, the Columbia Law and Economics Workshop, the UVA Law and Economics Workshop, the University of Chicago Law and Economics Workshop, the American Bar Association Groundbreakers Seminar, and the UC Berkeley Law and Economics Workshop.

1. Introduction

Web browsers and other search-facilitating applications allow users to input search queries, which are then directed to a search engine (for example, Google Search or Bing) to execute the search. But what determines which search engine is used in a given case? For many years, Google has paid most browsers and Android device makers to make Google Search the default search engine on their products.[1] Collectively, these deals are very broad in scope: about 70% of all search queries in the U.S. are made on browsers or other apps on which Google Search is the default search engine. Recently, these deals have become the focal point of *United States v. Google*, a landmark antitrust case that accuses Google of monopolizing the market for internet search.[2]

A user can switch the default search engine to a competing alternative (for example, Bing) through the settings menu. For a moderately tech-savvy user, this is not difficult. Thus, it is tempting to assume that default status is unlikely to have significant effects on consumer choice or competition. If this were true, default status would confer little value to a search engine. But Google pays tens of billions of dollars annually for its defaults. Indeed, in 2021 it paid about $20b to Apple alone. This immense price tag is a signal that defaults must be somewhat "sticky." That is, they must affect the search engine choice of a meaningful portion of consumers.

However, the competitive effects of these deals are not yet well understood. This paper attempts to fill that gap. Three important features of the search industry make this research problem especially rich. First, while we can usually focus on price effects to assess how a practice affects competition and consumer welfare, this is not possible here because search engines do not charge prices to consumers. Rather, they earn all revenue from advertisers. This relates to the second important feature, which is that search engines are two-sided media platforms: they deal simultaneously with two distinct but interdependent groups, consumers and advertisers. And defaults affect the two sides very differently. Finally, a search algorithm "learns" and improves with use.[3] This algorithmic learning process resembles network effects (Hagiu and Wright, 2023). This strongly influences the competitive impact of defaults.

I introduce a model of competition between two-sided search platforms, which are horizontally and vertically differentiated. One platform is dominant due to an exogenous quality advantage. Due to algorithmic learning, the number of consumers on a platform contributes to its overall quality. The platforms charge prices to advertisers, but they compete for consumers based on quality. I model this in a novel way. Each platform chooses a strategic variable called "ad intensity." When ad intensity rises, this increases the benefits advertisers get from running ads, but it also detracts from the user's search experience, reducing consumer utility. This framework generalizes numerous ideas in the literature about how media platforms exert market power

[1] Google pays Android device makers to preinstall its Chrome browser app, which uses Google Search by default. Because Android devices rarely preinstall multiple browsers, this makes Google Search the default search option on Android.

[2] The district court recently found Google liable. However, Google is expected to appeal the decision.

[3] Specifically, search engines use the data generated by their users' search activity to improve their algorithm. More users means more data and thus larger quality gains.

over consumers without charging them prices.[4]

When a platform acquires default status on a given browser, consumers on that browser must incur a switching cost to change to a different platform. But there is no cost to sticking with the default. This generates a biasing effect—it "nudges" consumers toward the default platform. It is analogous to imposing a tax on all but one of the products in a market.

In the "baseline model" (Section 2), I consider defaults that are maximally broad in the sense that they apply to all consumers in the market. As a result, at most one platform can have a default contract in place. This specification is useful for considering the case where a single search platform acquires default status on most or all browsers (as in the *Google* case).

So long as algorithmic learning effects are not too strong, a default agreement involving the dominant platform reduces consumer welfare. Like a tax, switching costs effectively make smaller rivals less attractive to consumers, leading some users to switch to the dominant platform. Due to algorithmic learning, this exodus of users reduces the quality of rivals' search algorithms, leading even more consumers to abandon them. This softens competition, allowing the dominant platform to raise ad intensity and advertising prices. But a default by the laggard makes competition more intense, because it counteracts the dominant platform's quality advantage to some extent. However, in the baseline model, the laggard's default can never raise welfare.[5] In fact, for most parameter values, it is more harmful than a default by the dominant platform, because it nudges the consumer base toward a less popular option.

Algorithmic learning magnifies the competitive effects of default agreements, because it makes consumer demand much more responsive to small changes in most relevant variables. This is significant, because it means that the switching costs introduced by a default need not be very large for the deal to have significant effects on competition. If algorithmic learning effects are extremely strong, then consumers could be better off under monopoly than competition.[6] In this case, the dominant platform's default will benefit consumers precisely because it reduces competition. However, empirical research suggests that algorithmic learning is subject to diminishing marginal returns. In an extension, I show that this makes it much less likely that the dominant platform's default could increase consumer welfare (Section 5.1). Diminishing returns also make it possible for a laggard's default to raise welfare.

Section 3 generalizes the baseline model so that a default applies to a single browser, with each browser accounting for some fraction of the consumer base. This allows for the possibility that all search platforms have defaults (on different browsers). This specification is useful for considering what the search industry would look like if default privileges were more evenly distributed across search platforms, rather than being concentrated in a single dominant firm. This leads to some qualitatively new effects. Most notably, defaults create spillovers—a default on one browser affects users on all other browsers. Interestingly, these spillovers can be positive,

[4]For example, the ad intensity level could represent the intrusiveness of ads on the search page, or the amount of personal data gathered on users. More abstractly, it could represent the amount of attention that consumers must "pay" to use a search platform.

[5]By analogy, decreasing a dominant firm's quality could make the market appear more competitive by reducing the firm's quality advantage, but this will generally reduce consumer welfare.

[6]This situation resembles natural monopoly and arises in many models of competition in network industries when network benefits are sufficiently large. See, for example, Nocke et al. (2007); Weyl and White (2014).

even if the default harms all users on the browser it applies to.

In this setting, the breadth of a default (the share of consumers directly affected by it) plays an important role in shaping its welfare effects. A laggard platform's default can raise aggregate consumer welfare, provided that it is sufficiently narrow. In this case, it is not algorithmic learning, but rather two other effects, that drive the welfare improvement. First, a laggard's default always generates net-positive spillovers, because it induces the dominant platform to behave more competitively across all browsers. Second, if the laggard's default is narrow, then the positive spillover affects many consumers, while the default's harmful effects are narrowly confined. By contrast, the dominant platform's default (even if narrow) harms consumers.

I also examine dynamic effects. Section 4 considers how defaults affect entry and investment by potential competitors. A default by the dominant platform induces potential entrants to invest less, reducing both the probability of entry and expected consumer welfare. By contrast, if a potential entrant (if successful) can obtain some default privileges, this increases its ex ante investment and raises the probability of entry, which can increase dynamic consumer welfare. Separately, in Section 5.3, I consider possible dynamic benefits of a dominant platform's defaults, such as promoting entry in the browser market. These possible dynamic efficiencies could potentially justify allowing a dominant search platform to maintain some defaults, provided they are not too broad.

Overall, the results paint a nuanced picture. On one hand, they suggest that broad defaults by a dominant search engine are likely to be harmful to consumers. Under most plausible conditions, the harm is both static and dynamic: these defaults not only leave consumers worse of in the short run, but also suppress long-run competition by chilling entry and investment by rivals. On the other hand, the results also show that defaults are not always bad. When reasonably modest in scope, defaults (particularly by laggards) can stimulate competition, both in the short run and in the long run.

Section 5 considers some extensions. Section 6 then concludes by discussing policy implications. One relevant question is what remedy or regulatory interventions (if any) would be most appropriate. I discuss several possibilities that have been proposed. One noteworthy example involves requiring browsers or devices to implement a "choice screen." This is a screen that requires the user to select her preferred search engine from a list. (A real-world example is shown in Appendix B.1.) Choice screens eliminate defaults.[7] However, this is unlikely to be the best policy solution, because limited defaults (particularly by laggards) can have important procompetitive effects, as noted above. This section also discusses why search engine defaults are unlikely to result in lower prices of mobile devices. Finally, I emphasize that this paper does not suggest that default-like arrangements in other industries are likely to be anticompetitive.

[7]Defaults apply in the absence of an active choice over search engines, but a choice screen forces the consumer to make an active choice.

1.1. *Related Literature*

There is relatively little economic scholarship on search engine defaults. Decarolis et al. (2025) studies the effects of recent implementations of choice screens (which eliminated Google's defaults) in Europe, Russia, and Turkey. I connect my model to some of their results in Section 5.2. The European 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 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.

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, although they do not focus specifically on the search market. But there are substantial differences between our papers. Their model does not contain algorithmic learning, which plays a crucial role in shaping virtually all of the effects of search engine defaults. Additionally, they consider homogeneous one-sided firms, whereas this paper considers differentiated two-sided platforms. Finally, they focus heavily on the process by which firms bid for default status, while I focus mainly on how defaults affect competition. Overall, our papers are complementary but distinct.

Algorithmic learning, which is a key component of this research problem, is an example of data-related processes that have recently received attention in industrial organization. For example, Hagiu and Wright (2023) consider "data-enabled learning"—processes by which a firm exploits its own user data to improve some measure of quality. And Agrawal et al. (2022) and Gans (2024) study related processes ("data feedback loops") in certain artificial intelligence (AI) contexts. De Corniere and Taylor (2025) provide a systematic overview of how firms gather, use, and trade data in ways that are relevant to competition.

Alcott et al. (2024) run a field experiment with desktop search users to study how defaults and choice screens affect consumer choice. They find that mistaken quality perceptions about Bing reinforce the effects of Google's defaults. As a result, they argue that choice screens are not likely to help stimulate competition. Vásquez Duque (2025) discusses other recent empirical studies that reach similar conclusions. I reach the same conclusion about choice screens, but for different reasons.

Several papers consider platform "self-preferencing" or "biased intermediation" (for example, De Corniere and Taylor, 2019; Hagiu et al., 2022; Padilla et al., 2022). In these papers, a platform monopolist also sells its own product, leading it to compete with other sellers on the platform. The platform can then favor its own product over those of competing sellers in some way.[8] Like these papers, mine surrounds a kind of biasing effect introduced by platforms. However, my paper focuses on how defaults affect competition *between platforms*, whereas self-preferencing papers focus on biasing of competition *between sellers on a platform*. My focus on search engines leads to substantial differences from papers on self-preferencing.[9]

[8]The manner in which the platform favors its own product varies among different papers. For example, in De Corniere and Taylor (2019), the platform recommends its own product more often, while in Hagiu et al. (2022) the platform may just remove rival sellers from the platform altogether.

[9]For example, algorithmic learning is a core feature of my model, but is not present in self-preferencing

Heidhues et al. (2024) explore the use of defaults as a self-preferencing strategy by digital ecosystems. In their model, an ecosystem makes its own product a default for all of its ecosystem users. The authors study how the potential to engage in self-preferencing leads ecosystems to acquire firms in adjacent markets, and how this pattern affects innovation and entry.

A number of theory papers focus on other aspects of search engine markets, such as auctions for ad slots on a search page (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 (Prat and Valletti, 2022). 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 some of its apps as a condition for gaining access to the Google Play store.

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. Due to algorithmic learning, a consumer on a given platform imposes positive externalities on all other consumers on the same platform. The model is of the well-known "competitive bottleneck" variety,[10] 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 is maximally broad—it affects all consumers equally, regardless of what browser they use. Hence, at most one platform can have default status. Section 3 generalizes the model so that a default applies to just one 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" (Anderson and Jullien, 2015) or "attention costs" (Newman, 2015; Gal and Rubinfeld, 2016).

papers. And my paper involves ad-supported media platforms, not transaction platforms that connect buyers and sellers.

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

- The platform could gather more data on its users. This is good for advertising but bad for consumers who value their privacy (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 (Garcia, 2023) or to enable advertising sellers to charge higher prices (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 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.

α_i represents the amount of unwanted advertising activity that consumers are subjected to when they use platform i. For example, it could represent the number or intrusiveness of ads on the search page; the amount of user data that the platform gathers; or the extent to which search results are distorted for advertisers' benefit. The ad intensity level can also be interpreted more abstractly as the amount of attention that consumers must "pay" to use the search platform.

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 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
$\widetilde{z}_i \equiv \alpha_i - p_i$	The marginal advertiser type on platform i
V_i	The exogenous quality of platform i's search technology
σ	The switching cost generated by a default contract
δ_i	A dummy with $\delta_i = 1$ iff platform i is the default
$\Delta \equiv V_1 - V_2 + (\delta_1 - \delta_2)\sigma$	Platform 1's technological advantage over platform 2
η	The strength of algorithmic learning effects
$\mu(\widetilde{z}_i) \equiv F(\widetilde{z}_i)/f(\widetilde{z}_i)$	The inverse semi-elasticity of advertising demand on platform i
$u_i(x)$	A consumer's utility from joining platform i
$\pi_i(z)$	An advertiser's payoff from joining 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 $\underline{z} \in \mathbb{R}$. It may or may not have a finite upper bound. Thus, the type space could be compact, $[\underline{z}, \overline{z}]$, or unbounded, $[\underline{z}, \infty)$.

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 $(\alpha_i - z)X_i$.[11] Here $X_i \in [0, 1]$ is consumer demand for platform i, which specifies the measure of consumers who participate on the platform. 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 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 $t \times \textit{distance}$, 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. Because platform 1 is dominant, we assume $V_1 > V_2$. This will ensure that it attracts a majority share of consumers in equilibrium ($X_1 > X_2$). The second term, $X_i\eta$, captures algorithmic learning benefits, with $\eta \geq 0$ giving the rate at which quality rises as a platform gains users. Finally, $h\alpha_i$ (where h is a positive scalar) is the disutility consumers experience due to unwanted advertising activities.[12] Intuitively, one can think of $h\alpha_i$ as the "price" a consumer pays to use platform i.

When one platform is the default, a consumer who wishes to use the other platform must incur

[11] This can be generalized so that the gross benefit is $(\gamma + \alpha_i - z)X_i$ for some fixed $\gamma \in \mathbb{R}$. However, because we do not restrict the sign of $\underline{z}$, the parameter γ can be normalized to zero without loss of generality.

[12] h captures how strongly consumers dislike platforms' advertising activities.

a switching cost of $\sigma \geq 0$. But there is no cost to using the default platform. Hence, the default generates a bias that favors the default platform. Define the dummy δ_i by

$$\delta_i \equiv \begin{cases} 1 & \text{if platform } i \text{ is the default} \\ 0 & \text{otherwise.} \end{cases}$$

Both platforms cannot be the default, so we impose $\delta_1 + \delta_2 \leq 1$. But we allow for the possibility that neither platform is the default: $\delta_1 = \delta_2 = 0$. We will associate this possibility with a choice screen, since a choice screen eliminates defaults.[13]

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

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

Notice that ad intensity (α_i) enters into consumer utility, but advertiser participation (Z_i) does not.[14] We assume the consumer side of the market is fully covered in all equilibria: $X_1 + X_2 = 1$. The marginal consumer location, defined implicitly by $u_1(\widetilde{x}) = u_2(\widetilde{x})$, is

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

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. A default by the dominant platform increases this advantage, while a default by the laggard reduces it.

By inspection of (6), the biasing effects of defaults manifest as vertical shifts in the demand functions $X_i = X_i(\alpha_1, \alpha_2)$. 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 cost σ is small.

Throughout the paper we assume that $t > \eta$. This is necessary for both platforms to be active in equilibrium.[15] For simplicity, we will also assume that σ is not large enough to make Δ negative when the laggard obtains default status (that is, we assume $\sigma < V_1 - V_2$). This ensures that platform 1 remains the market leader ($X_1 > X_2$) even when platform 2 is the

[13] See footnote 7. An example of a choice screen is depicted in Appendix B.1.

[14] This is because consumers care mainly about the amount of advertising activity they are subjected to when they use the search engine (for example, the number of ads *on the search page*), not the total number of advertisers who use the platform. Ad intensity is intended to capture the former. I explain this point further in Appendix B.4.

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

default.[16]

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 $\widetilde{z}_i$, where $\widetilde{z}_i \equiv \alpha_i - p_i$ is the marginal advertiser type on platform i (that is, the type such that $\pi_i(\widetilde{z}_i) = 0$). Notice that $\widetilde{z}_i$ is a good proxy for advertiser demand (Z_i), because they are directly related by $Z_i = mF(\widetilde{z}_i)$. Once $\widetilde{z}_i$ and p_i are chosen, the ad intensity level is pinned down as $\alpha_i = \widetilde{z}_i + p_i$.

Platform i's profits are given by

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

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

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

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

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

Evaluated at $\widetilde{z}_i$, this gives the inverse semi-elasticity of demand for advertising services on platform i.[17] As others have noted, the function $\mu(\cdot)$ is useful in characterizing demand (Johnson, 2017). Also, its equilibrium value is a good proxy for market power (Weyl, 2010).

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 $\mu'(z) > 0$.

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 (or neither). However, our second assumption stipulates that it cannot be *too* convex.

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

[16] It is straightforward to extend the results to the case where the laggard can overtake the dominant platform.

[17] To see this, note that

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

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.

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.[18] In Section 2.4 we will give an economic justification for assumption 2 based on revealed preference (see footnote 25). The final assumption simply ensures that each platform's optimal choice of $\widetilde{z}_i$ is always interior.

Assumption 3: $\widetilde{z}^{max} \equiv \mu^{-1}\left(\frac{2(t-\eta)}{h}\right)$ is in the support of F.

The marginal type $\widetilde{z}^{max}$ is what a monopolist would pick. In all equilibria with both platforms active, they want to pick values $\widetilde{z}_i \in (\underline{z}, \widetilde{z}^{max})$.[19] Assumption 3 ensures this is always possible.

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*

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

and

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

Equation (11) is simply the standard inverse elasticity pricing rule.[20] 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 ratio is equal to the ratio of consumer-side market shares:

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

Among other things, this implies that a platform's ability to charge high prices to advertisers is linked to its relative popularity among consumers (which depends on Δ). This reflects indirect network effects. As an illustration, consider the equilibrium when the model is linear:

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

[19] Recall that the support of F is assumed to be a (possibly unbounded) interval with a finite lower bound $\underline{z}$.

[20] Because $\mu(\widetilde{z}_i) = p_i/\varepsilon_i$ (see footnote 17), equation (11) can be rewritten as $\varepsilon_i^* = 1$, which is precisely the inverse elasticity pricing rule under costless production.

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*

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

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

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

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

In the linear model, $\mu(z)$ is simply the identity function, which is why $\widetilde{z}_i^* = p_i^*$ in this example.

2.4.1. Comparative Statics

When $\Delta = 0$, the equilibrium is fully symmetric.[21] The proposition below clarifies how changes in Δ (and, by extension, defaults) 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) *$\widetilde{z}_1^*$, p_1^*, α_1^*, and X_1^* are strictly increasing in Δ, while $\widetilde{z}_2^*$, 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) *$\widetilde{z}_i^*$, p_i^*, α_i^* $(i = 1, 2)$, and X_2^* are strictly decreasing in η, while X_1^* is strictly increasing in η.*

Part (i) tell us how default agreements affect market outcomes. When the dominant platform becomes the default it gains market power, allowing it to raise its ad price and its ad intensity level while also capturing more users on both sides. The laggard is weakened by the default. It mitigates the damage by reducing its price and ad intensity, but it still loses users on both sides. Market concentration thus increases on both sides. 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 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

[21] See the proof of Proposition 1.

is just one of several channels through which it affects consumer welfare. We consider all the different channels in the next subsection.

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

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.[22] 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.[23]

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.[24] As a result, defaults can potentially raise antitrust concerns even if it is not especially difficult for most consumers to switch search engines.

2.4.2. 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.[25]

Proposition 3. *Platform 1's willingness to pay for default status is strictly larger than that of platform 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.

[22] One consequence of this is that the threshhold Δ^{crit} from Proposition 1 is strictly decreasing in η.

[23] 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.

[24] 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). These shifts can be large even if σ is small.

[25] 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.

2.4.3. Consumer Welfare

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

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

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.[26] We want to know how a default contract affects consumer welfare relative to the benchmark where neither platform is the default (for example, 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[27]

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

Hence, the marginal effect on aggregate consumer welfare is just the share-weighted sum of the utility effects experienced by users on each platform.[28] 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. Additionally, because users of the non-default platform must switch away from the default, they also face a marginal switching cost of $\frac{\partial(-\sigma)}{\partial \sigma} = -1$.

Proposition 4. *Suppose platform i is the default, and let $j \neq i$.*

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

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

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

(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.[29] As for 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 5. *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 if* $\eta \leq \frac{2}{3}t$*. 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 so large that the gains to users on the dominant platform outweigh the harm to other users. In this parameter range the default increases consumer welfare *precisely because* it suppresses competition. When η is very large ($\eta \gg \frac{2}{3}t$), the situation resembles natural monopoly in that consumers are better off under monopoly than under competition. In that case, any anticompetitive arrangement that increases market concentration (for example, a merger) will raise consumer welfare. Importantly, however, if algorithmic learning benefits diminish with scale (as empirical work suggests), it is much less likely that the dominant platform's default could benefit consumers (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 *aggregate* algorithmic learning benefits, which are given by $X_1^*(X_1^*\eta) + X_2^*(X_2^*\eta) = HHI \times \eta$, where HHI is the Herfindahl index. Thus, for a default to increase consumer welfare, it must increase market concentration. But the laggard's default does just the opposite. However, as we will see in Sections 3-5, a laggard's default can raise welfare when we account for certain phenomena that are not present in the baseline game.

Suppose that one platform must be the default. Would consumers prefer it to be the laggard or the dominant platform? Proposition 5 implies that they prefer the latter option when η is very large. But what if η is small? Even in this case, consumers will *usually* prefer the default to favor the dominant platform. This is because the dominant platform is more popular ($X_1^* > X_2^*$), so its default nudges most consumers toward the same option they would have picked anyway.[30] 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.

[29] In fact, for these users, the negative switching cost effect always dominates the positive competition effect, so their utility would fall even without a negative learning effect.

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

Proposition 6. *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. Otherwise, the laggard's default is more harmful than the dominant platform's.*

Our assumptions do not require that $\mu''(\cdot)$ is negative, nor that η is small, so most allowed specifications of $\mu(\cdot)$ and η will not satisfy the condition in Proposition 6. 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.[31]

2.4.4. 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 7. *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.

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 defaults on most or all browsers. Although the results suggest that this is usually harmful, it does not follow that defaults are always bad. A common theme in the economics of vertical restraints is that, in small doses, such restraints are often procompetitive.

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

The antitrust concerns arise mainly when they are implemented on broad scales by a dominant firm.[32] Consistent with that trend, this section explains that the breadth of a platform's defaults is a major factor in determining their impact on consumer welfare.

We now generalize the model so that defaults have limited scope. Each default contract applies to a single browser, and thus directly affects only subset of consumers. Both platforms have defaults (on different browsers). Of course, if the allocation of defaults were determined entirely through competitive bidding, the dominant platform would acquire defaults on all browsers.[33] But we can view this generalization as capturing the situation where antitrust restricts the breadth of the dominant platform's defaults, similar to how antitrust treats other vertical restraints. This ensures that the laggard can acquire some default privileges, too.

This generalization leads to some qualitatively new effects. When defaults have limited breadth, they generate spillovers: a default on one browser affects users on other browsers. And, interestingly, these spillovers can be positive, even if the default harms all consumers on the browser it applies to. The breadth of a default thus becomes very important, as it determines not only how many consumers are directly affected by the default, but also how many consumers benefit from its positive spillovers.

There are now two browsers, "browser 1" and "browser 2." The fraction of consumers who use browser k is $\lambda^k \geq 0$, where $\lambda^1 + \lambda^2 = 1$. Platform 1 is the default on browser 1, while platform 2 is the default on browser 2. We interpret λ^i as the breadth of platform i's default. The switching costs associated with these defaults are denoted $\sigma^1 \geq 0$ and $\sigma^2 \geq 0$.[34] Throughout this section we assume that $\eta \leq \frac{2}{3}t$. It is in this parameter range that the interesting new results appear.[35] And we make the following additional assumptions:

- A consumer's browser selection is exogenously fixed and independent of her location, x.
- Each platform chooses a single ad intensity level α_i, which applies to both browsers.
- Algorithmic learning 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. 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 i is denoted $u_i^k(x)$. Thus, there are four utility functions, one for each browser-search platform

[32] For example, allowing video game consoles to acquire exclusive rights over a handful of video games is likely procompetitive (Lee, 2013). But if a dominant console attempted to acquire exclusive rights over *all* popular games, that would surely be anticompetitive and antitrust would prohibit it.

[33] This follows from the same argument underlying Proposition 3.

[34] In general $\sigma^1 \neq \sigma^2$, as it may be easier to switch search engines on one browser than the other. Also, a default contract may include restrictions that raise switching costs. For example, some of Google's deals prevent device makers from preinstalling rival search applications.

[35] When $\eta > \frac{2}{3}t$, the welfare results are qualitatively the same as in baseline game—the dominant platform's default can raise consumer welfare, but the laggard's cannot.

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.[36] 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. $\dfrac{\partial X_i^{j*}}{\partial \sigma^i} < 0 < \dfrac{\partial X_i^*}{\partial \sigma^i} < \dfrac{\partial X_i^{i*}}{\partial \sigma^i}$ *for each* i *and* $j \neq i$.[37]

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*. 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 $X_1^{2*} < X_2^{2*}$).

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

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

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

[37] The inequality $\partial X_i^{j*}/\partial\sigma^i < 0$ hinges on the assumption $\eta \leq \frac{2}{3}t$, but the others do not.

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. *In the model with browser-specific defaults:*

(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.[38] 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

	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

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

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 5 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 algorithmic learning 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.[39]

The masses of the different consumer groups also matter, as they determine how many users experience a given utility effect. These masses depend both on the platforms' browser-level market shares and on the breadth of their defaults (λ^1 and λ^2). Consider the masses in the top row of the tables. Because we know Google is more popular than Bing overall ($X_1^* > X_2^*$), Proposition 8 implies that $X_1^{1*} > X_1^* > \frac{1}{2} > X_2^* > X_2^{1*}$. It follows that $\lambda^1 X_1^{1*}$ is much larger than $\lambda^1 X_2^{1*}$, and hence Bing's default benefits far more Safari users than it harms. This explains why the laggard's default always generates net-positive spillovers.

However, turning to Google's default (Table 1), the relevant group masses (bottom row) are less likely to generate a net-positive spillover. 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 spillovers can be positive raises the question of whether a default could raise aggregate consumer welfare. The answer is yes.

Proposition 10. *In the model with browser-specific defaults:*

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

(ii) If λ^2 is sufficiently small, platform 2's default 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, in constrast to the baseline game, a default can raise consumer welfare even if algorithmic learning benefits are not terribly strong ($\eta \leq \frac{2}{3}t$). But there are two important caveats.

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

First, only the laggard's contract can potentially raise consumer welfare. Second, the laggard's contract can only raise welfare if it is sufficiently narrow in the sense that λ^2 is 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 the laggard's default is broad (large λ^2), we already know from the baseline game that this will reduce consumer welfare.

The proof gives a numerical example where the laggard's default raises aggregate consumer welfare. When demand is linear, this can happen only if σ^1 is positive. A positive value of σ^1 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 algorithmic learning effects (which detract from its consumer welfare benefits).

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 dynamic model to study how defaults affect entry and investment by potential entrants. A later section then considers some additional possible dynamic effects that have been widely discussed in connection with the *Google* case (Section 5.3).

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 model. 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 $\lim_{y \to 0} \rho'(y) = \infty$.[40] Finally, we impose $\eta \leq \frac{2}{3}t$, which ensures that entry is better for consumers than monopoly.

Although the dominant platform is always willing to outbid laggards for defaults, antitrust policy can lead to various allocations of defaults, which the platforms can anticipate ex ante. For example, if antitrust prohibits the dominant platform from obtaining broad defaults, then the entrant (if successful) will be able to acquire a default. Alternatively, antitrust could impose no restrictions on the dominant platform's defaults, or it could ban defaults altogether, forcing browsers to offer choice screens. We assume the operative policy is announced before the game starts. We need not model this policy announcement explicitly; we can simply assume that both platforms know ex ante which of them (if any) will end up with default status. For simplicity, we assume that defaults are maximally broad, as in the baseline model.[41]

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 probability of entry is $(1 - G(\widetilde{v}))\mathbb{E}[\rho(y^*(v))|v > \widetilde{v}]$. Defaults affect the equilibrium as follows:*

(i) If the platforms anticipate that the dominant platform will be 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) If the platforms anticipate that the potential entrant (if successful) will be 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, when an potential entrant (if successful) can obtain default status, the dynamic effects are procompetitive: the deal induces 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

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

[41] Thus, at most one firm can have default status. It is straightforward to extend this model so that both platforms can potentially obtain defaults on different browsers, as in the previous section.

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, algorithmic learning effects contribute an amount $X_i\eta$ to the perceived quality of platform i. This assumes that learning benefits increase linearly with scale. But, in practice, the quality benefits derived 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 learning benefits fall from $\overline{\eta}$ to $\underline{\eta}$, where $\overline{\eta} \geq \underline{\eta}$. In other words, algorithmic learning benefits now take the form

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

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.[42] 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 results given in this section hold 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.[43] The parameter $r \in [0,1]$ governs the relative magnitudes of $\overline{\eta}$ and $\underline{\eta}$, and thus captures the degree to which algorithmic learning benefits diminish with scale. Notice that $r = 0$ just results in the baseline model. This

[42] 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 $\hat{X}\overline{\eta}$).

[43] As in the baseline game, we require $\eta < t$ for both firms to be active. However, $\overline{\eta}$ and $\underline{\eta}$ need not satisfy that constraint individually.

parameterization is helpful, because equilibrium strategies and demand levels depend only on 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. The first main result of this section is that, when there are diminishing returns to data, default agreements by a dominant platform are more detrimental to consumers.

Proposition 12. *Suppose platform 1 is the default. Under diminishing returns, the default's impact on aggregate consumer welfare is strictly decreasing in r.*

This result makes intuitive sense. When there are diminishing returns to data, the dominant platform's default increases its own quality only slightly, but it reduces the laggard's quality significantly. This asymmetry makes the default more harmful. It also makes it much less likely that the dominant platform's default could raise aggregate consumer welfare, as that possibility requires that the default would greatly increase the dominant platform's quality.

Turning this argument on its head, diminishing returns make it possible for the laggard's default to raise aggregate consumer welfare, in contrast to the baseline model.

Proposition 13. *Suppose platform 2 is the default. Under diminishing returns, the default increases aggregate consumer welfare whenever the difference $\overline{\eta} - \underline{\eta}$ is sufficiently large.*

The welfare improvement arises when algorithmic learning effects are initially strong (large $\overline{\eta}$) but drop off significantly (small $\underline{\eta}$). In this case, the laggard's default greatly increases its own quality, but only slightly reduces the dominant platform's quality.

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.[44] Do defaults by the two platforms elicit 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.

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: $\mathcal{S}_2 > \mathcal{S}_1$. This is supported by empirical evidence on the effects of choice screen mandates (which eliminated Google's default status on Android) across different foreign jurisdictions in which Google faces varying degrees of competition.[45]

By contrast, in the baseline model, the relationship between $\mathcal{S}_1$ and $\mathcal{S}_2$ is ambiguous in general. It depends on details about the curvature of demand.

Proposition 14. *In the baseline model:*

(i) $\mathcal{S}_1 = \mathcal{S}_2$ *if demand is linear (that is, 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, both demand functions might shift by a larger amount when the laggard is the default. However, in the baseline model, 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 effects are likewise symmetric ($\mathcal{S}_1 = \mathcal{S}_2$). 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.[46] 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 the two types differ only in the size of

[44] 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 $\mathcal{S}_2 \equiv X_2^*(V_1 - V_2 - \sigma) - X_2^*(V_1 - V_2)$.

[45] For example, Decarolis et al. (2025) 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 European 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%.

[46] 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.

their switching costs, a consumer's type doesn't matter when a choice screen is implemented, since a choice screen makes switching costs irrelevant.[47]

In this extension, the magnitudes of the demand shifts are no longer symmetric.[48] Specifically, a default by the laggard generates larger shifts than a default by the dominant platform.[49] 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 15. *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). In other words, as the dominant firm's defaults cause more and more harm over time, they start to appear increasingly feeble at the margin. Accordingly, if evidence suggests that eliminating the default would have only minor effects on market shares in the short run, this does not imply that the defaults are not harmful.

5.3. 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 like Apple from creating its own search engine? This section briefly considers what the model implies about the plausibility of these possibilities.

Proposition 16. *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.

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

[48] See the proof of Proposition 15 for details.

[49] 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.

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

These results shed light on the dynamic questions mentioned above.

Promoting competition in the browser market. Mozilla (creator of the Firefox browser) 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 3 implies the dominant platform is willing to pay more than the laggard. Second, from part (i) of Proposition 16, 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.[50] 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, because they only apply to a subset of Google's defaults, it is not clear why these arguments would justify Google in maintaining such a broad set of defaults. (Collectively, Google's defaults cover ∼70% of all domestic search queries.) This is underscored by the fact that broader defaults tend to be more harmful to static consumer welfare.[51] Accordingly, these possible dynamic efficiencies may caution against a categorical ban on defaults, but they do not imply that there should be no antitrust limits at all.

[50] Even if platforms are highly 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.

[51] This follows from arguments in Section 3.

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 whose effects on competition and consumer welfare are not yet well understood. This paper attempts to fill that gap. The results suggest that, under plausible conditions, broad defaults by a dominant search platform are likely to be anticompetitive. Aside from reducing static consumer welfare, they are likely to suppress entry and investment by laggard competitors.

Algorithmic learning magnifies the competitive effects of defaults. As a result, defaults can have a significant effect on competition even if switching costs are not very large. If algorithmic learning effects are sufficiently strong, consumers might prefer monopoly over competition. If so, Google Search might be something like a natural monopoly. In such a world, a dominant platform's default agreements would benefit consumers *precisely because* they suppress competition and shift the market structure closer to monopoly. However, most research suggests there are diminishing returns to user data. As explained in Section 5.1, this makes it much less likely that consumers might benefit from broad defaults by a dominant platform.

However, the results also indicate that defaults (particularly by laggards) are sometimes procompetitive. A narrow default agreement by a laggard can raise consumer welfare by forcing the dominant platform to behave more competitively. Diminishing returns to user data also make it more likely that a laggard's default could raise welfare. And allowing laggards and entrants to obtain default placements can stimulate entry and investment. This fits neatly with other research suggesting that vertical restraints may play an important entry-facilitating role in network industries.[52] Indeed, even defaults by a dominant platform like Google may have some dynamic benefits, such as by promoting entry in the browser market (Section 5.3.) Importantly, however, it is not clear why those possible dynamic benefits would require Google to maintain such a broad set of defaults.

These findings have important implications for antitrust policy toward defaults. On one hand, they cast doubt on the wisdom of any policy that would elimiante default altogether. In some situations, defaults stimulate competition or entry, and we don't want to eliminate those procompetitive benefits. On the other hand, the results do suggest that some antitrust limits are appropriate. Broad defaults by a dominant platform are likely to undermine both static and dynamic competition. And the potential procompetitive justifications for defaults would not seem to require such broad defaults. These points suggest a policy of treating defaults in roughly the same way antitrust treats exclusive dealing: lawful in moderation, but not in excess. Under such a policy, a search platform's defaults would potentially arouse antitrust concerns only if they become sufficiently broad in scope.

Remedy considerations. If a court finds antitrust liability, what is the right remedy? A detailed

[52]For example, Lee (2013) finds that exclusive rights over video games helped to facilitate successful entry by laggard video game consoles. 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.

discussion of this question is beyond the scope of this paper but I will offer a few comments. The simplest remedy would be an injunction that terminates some or all of the Google's contracts. That would leave device makers and browsers free to enter into similar agreements with smaller search engines like Bing. The court might also require Google to give up its default on its own Chrome browser (at least temporarily), perhaps by offering users a choice screen.[53]

A court could also order Google to share its user data, allowing smaller rivals to draw algorithmic learning benefits from Google's data. Hagiu and Wright (2023) finds that this would raise consumer welfare if the market leader is sufficiently far ahead of the laggard. However, Martens (2024) notes that data sharing could inadvertently reduce aggregate algorithmic learning benefits if the dominant search platform shares data with multiple small laggards. This could be avoided by making data-sharing reciprocal rather than unidirectional. Some observers have also raised concerns that data sharing could raise concerns about user privacy.

Mandatory choice screens are another widely-discussed possibility, although this is probably not on the table in the *Google* case.[54] Such a mandate would eliminate defaults completely. However, the results suggest that this is unlikely to be a desirable solution, as defaults can often stimulate competition, at least if they are not too broad. For example, a choice screen could inadvertently harm competition by taking away a valuable entry-facilitating tool.

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. (2025) studied the effects of various choice screens implemented overseas. They found that the resulting elimination of Google's defaults had no significant effect on device sales, suggesting there were no meaningful price effects. Additionally, the way payments are structured in default agreements is not conducive to significant price effects, because most device sales do not affect the payments received by the device maker. In Appendix B.3. I discuss these and other difficulties facing the subsidization argument.

Large payments as an entry deterrent? As noted in Section 5.3, 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. Ordinarily, when a threat of competitive entry induces an incumbent to offer its trading partners

[53]The Justice Department has requested that the court go farther than this and order Google to divest Chrome (and possibly Android as well). This strikes me as excessive and premature at this stage. The case for a breakup is strongest when a defendant's misconduct is unilateral and hard for a court to police, but Google's broad default coverage is driven mainly by contracts, which the court can simply terminate. While Google could still give itself the Chrome default unilaterally, the court could presumably prohibit that without requiring a divesture, as noted above. Thus, a breakup risks destroying significant economic value unnecessarily.

[54]The *Google* court's remedy order will only bind Google; it cannot force other firms like Apple to implement choice screens.

more favorable terms, we view this as a good thing, even if it means that no entry ultimately occurs. It is not clear why we should stray from that view here.

Default-like arrangements in other settings. 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.[55] 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 manufacturing costs in any relevant market. Also, the overwhelming majority of such restrictions are not implemented on broad scale contemplated in much of this paper.[56]

References

Agrawal, A., J. Gans, and A. Goldfarb (2022). *Prediction Machines, Updated and Expanded: The Simple Economics of Artificial Intelligence*. Harvard Business Press.

Alcott, H., J. C. Castillo, M. Gentzkow, L. Musolff, and T. Salz (2024). *Sources of Market Power in Web Search: Evidence from a Field Experiment*. Working paper.

Anderson, S. P. and B. Jullien (2015). The advertising-financed business model in two-sided media markets. In *Handbook of media economics*, Volume 1, pp. 41–90. Elsevier.

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

Armstrong, M. and J. Wright (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory 32*, 353–380.

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

Bet, G., R. D. Blair, and J. D. Donna (2022). The economic rationale of united states v. google. *The Antitrust Bulletin 67*(1), 23–39.

Chen, Y. and M. Schwartz (2024). Assigning default position for digital goods: Competition, regulation, and welfare. Technical report.

Choi, J. P., D.-S. Jeon, and M. Whinston (2023). Tying in markets with network effects. *working paper*.

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

[55] 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.

[56] Google's defaults cover 70% of the market—something that only a highly dominant firm could accomplish.

De Corniere, A. and G. Taylor (2019). A model of biased intermediation. *The RAND Journal of Economics 50*(4), 854–882.

De Corniere, A. and G. Taylor (2025). Data and competitition: A simple framework. *The RAND Journal of Economics (forthcoming)*.

Decarolis, F., M. Li, and F. Paternollo (2025). Competition and defaults in online search. *American Economic Journal: Microeconomics (forthcoming)*.

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

Gans, J. S. (2024). *Market Power in Artificial Intelligence*. NBER Working Paper.

Garcia, D. (2023). Search engine competition.

Gilbert, R. J. and D. M. Newbery (1982). Preemptive patenting and the persistence of monopoly. *The American Economic Review*, 514–526.

Hagiu, A., T.-H. Teh, and J. Wright (2022). Should platforms be allowed to sell on their own marketplaces? *The RAND Journal of Economics 53*(2), 297–327.

Hagiu, A. and J. Wright (2023). Data-enabled learning, network effects, and competitive advantage. *The RAND Journal of Economics 54*(4), 638–667.

Heidhues, P., M. Kösters, and B. Kîoszegi (2024). A theory of dgital ecosystems. Technical report, ECONtribute Discussion Paper.

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

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

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

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

Nocke, V., M. Peitz, and K. Stahl (2007). Platform ownership. *Journal of the European Economic Association 5*(6), 1130–1160.

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

Padilla, J., J. Perkins, and S. Piccolo (2022). Self-preferencing in markets with vertically integrated gatekeeper platforms. *The Journal of Industrial Economics 70*(2), 371–395.

Peitz, M. and M. Reisinger (2015). The economics of internet media. In *Handbook of media economics*, Volume 1, pp. 445–530. Elsevier.

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

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

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

Vásquez Duque, O. (2025). The magical number 2 (minus two): An empirical analysis on the efficacy of choice screens to increase competition in digital markets. *Mimeograph*.

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

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

Appendix A: Proofs

Proof of Proposition 1

Proof. Platform i's profits are $\Pi_i = mF(\widetilde{z}_i)X_i p_i$. Viewing this as a function of $(\widetilde{z}_i, p_i)$, 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 $\widetilde{z}_i$ and p_i are always interior (that is, they satisfy $F(\widetilde{z}_i) < 1$ 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 $\widetilde{z}_i$ is

$$f(\widetilde{z}_i)X_i = \frac{h}{2(t-\eta)}F(\widetilde{z}_i) \quad \Longleftrightarrow \quad X_i = \frac{h}{2(t-\eta)}\mu(\widetilde{z}_i), \tag{28}$$

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

$$X_i = \frac{h}{2(t-\eta)}p_i. \tag{29}$$

Combining these yields $p_i = \mu(\widetilde{z}_i)$, 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 (29) and solving for prices, we get

$$p_1^{br} = \frac{t-\eta+\Delta-h(\widetilde{z}_1-\widetilde{z}_2)+hp_2}{2h}, \quad p_2^{br} = \frac{t-\eta-\Delta+h(\widetilde{z}_1-\widetilde{z}_2)+hp_1}{2h}. \tag{30}$$

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

$$p_1^\dagger = \frac{3(t-\eta) + \Delta - h(\widetilde{z}_1 - \widetilde{z}_2)}{3h}, \qquad p_2^\dagger = \frac{3(t-\eta) - \Delta + h(\widetilde{z}_1 - \widetilde{z}_2)}{3h}. \tag{31}$$

Suppose that the equilibrium choices of $\widetilde{z}_i$ satisfy $|\widetilde{z}_1 - \widetilde{z}_2| \to 0$ 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 $p_i^\dagger = \mu(\widetilde{z}_i)$ in equilibrium. Plugging this into (31) and rearranging, we get the following pair of equations:

$$\mu(\widetilde{z}_1) + \frac{\widetilde{z}_1}{3} = \frac{t-\eta}{h} + \frac{\Delta}{3h} + \frac{\widetilde{z}_2}{3}, \qquad \mu(\widetilde{z}_2) + \frac{\widetilde{z}_2}{3} = \frac{t-\eta}{h} - \frac{\Delta}{3h} + \frac{\widetilde{z}_1}{3}. \tag{32}$$

These equations implicitly define best response functions for $\widetilde{z}_1$ and $\widetilde{z}_2$. Implicit differentiation of the first equation w.r.t. $\widetilde{z}_2$ shows that the slope of platform 1's best response function is $1/(1+3\mu'(\widetilde{z}_1)) \in (0,1)$. A symmetric argument applies to platform 2. Since the slopes of both reaction functions lie in $(0,1)$, there exists at most one pair $(\widetilde{z}_1, \widetilde{z}_2)$ satisfying both equations in (32). When $\Delta = 0$, both equations are satisfied by setting $\widetilde{z}_1 = \widetilde{z}_2 = \widetilde{z}^0$, where $\widetilde{z}^0$ is defined implicitly by $\mu(\widetilde{z}^0) = \frac{t-\eta}{h}$. This is the symmetric equilibrium of the game. This equilibrium has demand levels $X_1 = X_2 = \frac{1}{2}$ and $Z_1 = Z_2 = mF(\widetilde{z}^0) > 0$, 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 $\mu(\widetilde{z}_1^*) + \mu(\widetilde{z}_2^*) = \frac{2(t-\eta)}{h}$. As we will verify in the proof of Proposition 2, $\widetilde{z}_1^*$ is strictly increasing in Δ and $\widetilde{z}_2^*$ is strictly decreasing. (This implies that $\widetilde{z}_1^* > \widetilde{z}_2^*$, since they are equal when $\Delta = 0$.) Thus, both firms are active so long as $\widetilde{z}_2^* > \underline{z}$. Δ^{crit} is the value of Δ at which $\widetilde{z}_2^* = \underline{z}$, in which case we must have $\mu(\widetilde{z}_1^*) = 2(t-\eta)/h \Leftrightarrow \widetilde{z}_1^* = \widetilde{z}^{max} \equiv \mu^{-1}\big(2(t-\eta)/h\big)$. (Assumption 3 ensures this is feasible.) The corresponding prices are $p_1^* = \mu(\widetilde{z}^{max}) = 2(t-\eta)/h$ and $p_2^* = \mu(\underline{z}) = 0$. Plugging these strategy values into (31) we find

$$\Delta^{crit} = 3(t-\eta) + h\mu^{-1}\left(\frac{2(t-\eta)}{h}\right) - h\underline{z}. \tag{33}$$

□

Proof of Proposition 2

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

$$p_1^* = \frac{3(t-\eta) + \Delta - h(\widetilde{z}_1^* - \widetilde{z}_2^*)}{3h}, \quad p_2^* = \frac{3(t-\eta) - \Delta + h(\widetilde{z}_1^* - \widetilde{z}_2^*)}{3h}. \tag{34}$$

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

$$X_1^* = \frac{1}{2} + \frac{\Delta - h(\widetilde{z}_1^* - \widetilde{z}_2^*)}{6(t-\eta)}, \qquad X_2^* = \frac{1}{2} - \frac{\Delta - h(\widetilde{z}_1^* - \widetilde{z}_2^*)}{6(t-\eta)}. \tag{35}$$

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

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

Let $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$, and let ∂_Δ be the derivative operator $\partial_\Delta = \frac{\partial}{\partial \Delta}$. Differentiating the equations in (36) w.r.t. Δ and rearranging, we find

$$\partial_\Delta \widetilde{z}_1^* = \frac{1 + h\partial_\Delta \widetilde{z}_2^*}{h(1+3\mu_1')}, \qquad \partial_\Delta \widetilde{z}_2^* = \frac{-1 + h\partial_\Delta \widetilde{z}_1^*}{h(1+3\mu_2')}. \tag{37}$$

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

$$\partial_\Delta \widetilde{z}_1^* = \frac{\mu_2'}{hM} > 0, \qquad \partial_\Delta \widetilde{z}_2^* = -\frac{\mu_1'}{hM} < 0, \tag{38}$$

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

$$\partial_\Delta p_1^* = \mu_1' \partial_\Delta \widetilde{z}_1^* = \frac{\mu_1'\mu_2'}{hM} > 0, \qquad \partial_\Delta p_2^* = \mu_2' \partial_\Delta \widetilde{z}_2^* = -\frac{\mu_1'\mu_2'}{hM} < 0 \tag{39}$$

and

$$\partial_\Delta X_1^* = \frac{h}{2(t-\eta)}\mu_1' \partial_\Delta \widetilde{z}_1^* = \frac{\mu_1'\mu_2'}{2(t-\eta)M} > 0, \qquad \partial_\Delta X_2^* = \frac{h}{2(t-\eta)}\mu_2' \partial_\Delta \widetilde{z}_2^* = -\frac{\mu_1'\mu_2'}{2(t-\eta)M} < 0. \tag{40}$$

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

$$\partial_\Delta \alpha_1^* = \frac{\mu_2' + \mu_1'\mu_2'}{hM} > 0, \qquad \partial_\Delta \alpha_2^* = -\frac{\mu_1' + \mu_1'\mu_2'}{hM} < 0. \tag{41}$$

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

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

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

Part (ii). For $\widetilde{z}_i^*$, 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 (36), 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 $\partial_\eta \widetilde{z}_i^*$ for $i = 1, 2$). Solving this system of

equations gives us

$$\partial_\eta \widetilde{z}_1^* = -\frac{3\mu_2' + 2}{hM} < 0, \qquad \partial_\eta \widetilde{z}_2^* = -\frac{3\mu_1' + 2}{hM} < 0. \tag{42}$$

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

$$\partial_\eta p_1^* = \mu_1' \partial_\eta \widetilde{z}_1^* = -\frac{3\mu_1'\mu_2' + 2\mu_1'}{hM} < 0, \qquad \partial_\eta p_2^* = \mu_2' \partial_\eta \widetilde{z}_2^* = -\frac{3\mu_1'\mu_2' + 2\mu_2'}{hM} < 0. \tag{43}$$

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 $\mu(\widetilde{z}_1^*) + \mu(\widetilde{z}_2^*) = \frac{2(t-\eta)}{h}$ into (12), we get the following expression for consumer-side market shares:

$$X_i^* = \frac{\mu(\widetilde{z}_i^*)}{\mu(\widetilde{z}_i^*) + \mu(\widetilde{z}_j^*)} \tag{44}$$

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:

$$\partial_\eta X_1^* = \frac{\mu_1'\mu_2'}{(\mu(\widetilde{z}_1^*) + \mu(\widetilde{z}_2^*))^2 hM} \left\{ 3(\mu(\widetilde{z}_1^*) - \mu(\widetilde{z}_1^*)) + 2\left(\frac{\mu(\widetilde{z}_1^*)}{\mu_1'} - \frac{\mu(\widetilde{z}_2^*)}{\mu_2'} \right) \right\}. \tag{45}$$

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

Proof of Proposition 3

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

$$\begin{aligned}
\partial_\Delta \sum_i \Pi_i^* &= \frac{mh}{2(t-\eta)} \sum_i \left\{ f(\widetilde{z}_i^*) \mu(\widetilde{z}_i^*)^2 + 2F(\widetilde{z}_i^*) \mu(\widetilde{z}_i^*) \mu_i' \right\} \partial_\Delta \widetilde{z}_i^* \\
&= \frac{mh}{2(t-\eta)} \sum_i \left\{ \mu(\widetilde{z}_i^*) + 2\mu(\widetilde{z}_i^*) \mu_i' \right\} F(\widetilde{z}_i^*) \partial_\Delta \widetilde{z}_i^* \qquad (\text{using } f = \tfrac{F}{\mu}) \\
&= \frac{m}{2(t-\eta)M} \left\{ \left(\mu(\widetilde{z}_1^*) + 2\mu(\widetilde{z}_1^*)\mu_1' \right) F(\widetilde{z}_1^*) \mu_2' - \left(\mu(\widetilde{z}_2^*) + 2\mu(\widetilde{z}_2^*)\mu_2' \right) F(\widetilde{z}_2^*) \mu_1' \right\} \qquad (\text{using (38)}) \\
&= \frac{m\mu_1'\mu_2'}{2(t-\eta)M} \left\{ \left(\frac{\mu(\widetilde{z}_1^*)}{\mu_1'} + 2\mu(\widetilde{z}_1^*) \right) F(\widetilde{z}_1^*) - \left(\frac{\mu(\widetilde{z}_2^*)}{\mu_2'} + 2\mu(\widetilde{z}_2^*) \right) F(\widetilde{z}_2^*) \right\}.
\end{aligned}$$

This is positive, because $\widetilde{z}_1^* > \widetilde{z}_2^*$ 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 4

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

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

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

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

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 5

Proof. As in earlier proofs, let $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. A default for platform i is equivalent to reducing V_j by σ. The derivatives of $\mathcal{W}_c$ are

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

where $j \neq i$. It is easy to verify that

$$\frac{\partial u_i^*}{\partial V_i} = 1 - \frac{\mu_j' + \mu_1'\mu_2'}{M} + \frac{\eta\mu_1'\mu_2'}{2(t-\eta)M}$$

is positive in all cases. Additionally, some algebra shows that the cross-derivatives

$$\frac{\partial u_j^*}{\partial V_i} = \frac{\mu_i' + \mu_1'\mu_2'}{M} - \frac{\eta\mu_1'\mu_2'}{2(t-\eta)M}$$

are positive whenever $\eta \leq \frac{2}{3}t$. Thus, $\eta \leq \frac{2}{3}t$ 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. Now suppose η is very large. Using the above equations, we find

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

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 6

Proof. A default for platform i is equivalent to reducing V_j by σ. Thus, 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 $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$, $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Using results from the proof of Proposition 2, we find

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

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 $\frac{\partial \mathcal{W}_c}{\partial V_1} \geq \frac{\partial \mathcal{W}_c}{\partial V_2}$. 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 7

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

$$\begin{aligned}\mathcal{W}_a &= \int_{\underline{z}}^{\widetilde{z}_1^*}(\widetilde{z}_1^* - z)X_1^* f(z)dz + \int_{\underline{z}}^{\widetilde{z}_2^*}(\widetilde{z}_2^* - z)X_2^* f(z)dz \\ &= \widetilde{z}_1^* X_1^* F(\widetilde{z}_1^*) - X_1^* \int_{\underline{z}}^{\widetilde{z}_1^*} z f(z)dz + \widetilde{z}_2^* X_2^* F(\widetilde{z}_2^*) - X_2^* \int_{\underline{z}}^{\widetilde{z}_2^*} z f(z)dz. \end{aligned} \tag{46}$$

Applying integration by parts to these integrals:

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

Plugging this into (46) and simplifying, we get

$$\mathcal{W}_a = X_1^* \int_{\underline{z}}^{\widetilde{z}_1^*} F(z)dz + X_2^* \int_{\underline{z}}^{\widetilde{z}_2^*} F(z)dz. \tag{47}$$

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 8

Proof. As in previous proofs, let $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$ 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

$$X_i^i = X_i + \frac{\lambda^j(\sigma^1 + \sigma^2)}{2t}, \qquad X_i^j = X_i - \frac{\lambda^i(\sigma^1 + \sigma^2)}{2t}. \tag{48}$$

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

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

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 $\frac{\partial X_i^{j*}}{\partial \sigma^i} < 0$. □

Proof of Proposition 9

Proof. As in previous proofs, let $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$ 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 4, 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:

$$\frac{\partial u_i^i}{\partial \sigma^i} = -\lambda^i \psi_i \qquad \frac{\partial u_j^i}{\partial \sigma^i} = -(1 - \lambda^i \psi_j) \tag{49}$$

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

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:

$$\begin{aligned} X_1^* \psi_1 - X_2^* \psi_2 &= X_1^* \frac{\mu_2' + \mu_1' \mu_2'}{M} - X_2^* \frac{\mu_1' + \mu_1' \mu_2'}{M} - (X_1^* - X_2^*) \frac{\eta \mu_1' \mu_2'}{2(t-\eta)M} \\ &= \frac{\mu_1' \mu_2'}{M} \left(\frac{X_1^*}{\mu_1'} - \frac{X_2^*}{\mu_2'} + (X_1^* - X_2^*) \left[\frac{2t - 3\eta}{2(t-\eta)} \right] \right) \\ &= \frac{h \mu_1' \mu_2'}{2(t-\eta)M} \left(\frac{\mu(\widetilde{z}_1^*)}{\mu_1'} - \frac{\mu(\widetilde{z}_2^*)}{\mu_2'} + (\mu(\widetilde{z}_1^*) - \mu(\widetilde{z}_2^*)) \left[\frac{2t - 3\eta}{2(t-\eta)} \right] \right) \end{aligned} \tag{51}$$

where the third line follows from (12). Assumptions 1 and 2 ensure that $\mu(\widetilde{z}_1^*) > \mu(\widetilde{z}_2^*)$ and $\frac{\mu(\widetilde{z}_1^*)}{\mu_1'} > \frac{\mu(\widetilde{z}_2^*)}{\mu_2'}$. Thus, the expression in (51) 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

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

where the final inequality follows from (51). This establishes part (ii). As for platform 1's spillover, it is

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

where the second equality uses (48). 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 $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$ 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

$$\frac{\partial \mathcal{W}_c}{\partial \sigma^2} = \lambda^2\Big[-X_1^{2*} + X_1^*\psi_1 - X_2^*\psi_2\Big]. \tag{52}$$

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 (48), 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

$$\frac{\partial \mathcal{W}_c}{\partial \sigma^2} = \lambda^2\Big[-X_1^{2*} + X_1^*\psi_1 - X_2^*\psi_2\Big] = \frac{\lambda^2}{18t}\Big[8\lambda^1\sigma^1 + (1+8\lambda^1)\sigma^2 - 9t - \Delta^0\Big]. \tag{53}$$

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 (that is, 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 (52) 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

$$V_1 - \widetilde{v} = \Delta^{crit}. \tag{54}$$

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

$$\rho'(y^*)\Pi_2^*(V_1, v) - 1 = 0. \tag{55}$$

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

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

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 $\mathcal{W}_c^M = \frac{1}{2}t$.[57] 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 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 $\widetilde{v}$ to $\widetilde{v}^- \equiv V_1 - \sigma - \Delta^{crit}$, it will necessarily raise the welfare contributions from all types $v \in (\widetilde{v}^-, \widetilde{v} + \varepsilon)$ for sufficiently small $\varepsilon > 0$. All of these types contribute positive amounts to expected consumer welfare. But absent the default, types $v \in (\widetilde{v}^-, \widetilde{v}]$ would contribute nothing to welfare and types $v \in (\widetilde{v}, \widetilde{v} + \varepsilon)$ would contribute very little, since $y^*(v)$ would be close to zero in a right-neighborhood of $v = \widetilde{v}$. Now suppose that G assigns almost all probability to types in $v \in (\widetilde{v}^-, \widetilde{v} + \varepsilon)$. 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. Let $\mathcal{W}_c(r)$ denote aggregate consumer welfare as a function of r. Recall from the proof of Proposition 6 that $\frac{\partial \mathcal{W}_c(r)}{\partial \sigma} = X_1^* \frac{\partial u_1^*}{\partial \sigma} + X_2^* \frac{\partial u_2^*}{\partial \sigma}$. As in previous proofs, let $\mu_i' \equiv \mu'(\widetilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Using the results from the proof of Proposition 2, we have

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

[57] When platform 1 is a monopolist, it sets α_1 so that $u_1(1) = 0$. That is, 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).

And, by a similar process, we find

$$\frac{\partial u_2^*}{\partial \sigma} = -\left(1 - \frac{\mu_1' + \mu_1'\mu_2'}{M} + \frac{\eta\mu_1'\mu_2'}{2(t-\eta)M}\right) - \frac{r\eta\mu_1'\mu_2'}{2(t-\eta)M} \tag{57}$$

Therefore

$$\frac{\partial \mathcal{W}_c(r)}{\partial \sigma} = -X_1^*\left(\frac{\mu_2' + \mu_1'\mu_2'}{M}\right) - X_2^*\left(1 - \frac{\mu_1' + \mu_1'\mu_2'}{M}\right) \tag{58}$$

$$+ (X_1^* - X_2^*)\frac{\eta\mu_1'\mu_2'}{2(t-\eta)M} - (X_1^* + X_2^*)\frac{r\eta\mu_1'\mu_2'}{2(t-\eta)M} \tag{59}$$

$$= \left.\frac{\partial \mathcal{W}_c(r)}{\partial \sigma}\right|_{r=0} - \frac{r\eta\mu_1'\mu_2'}{2(t-\eta)M} \tag{60}$$

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. This implies that $\frac{\partial \mathcal{W}_c(r)}{\partial \sigma}$, and by extension the default's cumulative welfare effect, is strictly decreasing in r. □

Proof of Proposition 13

Proof. Suppose the laggard is the default. Note that, because we assume $\sigma < V_1 - V_2$, there is no possibility that platform 2 overtakes platform 1 as the market leader. As before, let $\mu_i' \equiv \mu'(\tilde{z}_i^*)$ and $M \equiv \mu_1' + \mu_2' + 3\mu_1'\mu_2'$. Using the proof of Proposition 2, the marginal utility effects are

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

and

$$\frac{\partial u_2^*}{\partial \sigma} = -\frac{\partial X_2^*}{\partial \Delta}\overline{\eta} + h\frac{\partial \alpha_2^*}{\partial \Delta} = \frac{\mu_1'\mu_2'\overline{\eta}}{2(t-\eta)M} - \frac{\mu_1' + \mu_1'\mu_2'}{M}.$$

Recall from the proof of Proposition 6 that $\frac{\partial \mathcal{W}_c}{\partial \sigma} = X_1^*\frac{\partial u_1^*}{\partial \sigma} + X_2^*\frac{\partial u_2^*}{\partial \sigma}$. Thus, we have

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

After simplification:

$$\frac{\partial \mathcal{W}_c}{\partial \sigma} = \underbrace{X_1^*\left(-1 + \frac{\mu_2' + \mu_1'\mu_2'}{M}\right)}_{\in(-1,0)} + \underbrace{X_2^*\left(-\frac{\mu_1' + \mu_1'\mu_2'}{M}\right)}_{\in(-1,0)} + \frac{\mu_1'\mu_2'\eta}{2(t-\eta)}\Big[r - (X_1^* - X_2^*)\Big].$$

Focus on the third term. Note that $\overline{\eta} - \underline{\eta} = 2\eta r$, so this difference is large when η and r are large. Recalling that $0 \le \eta < t$ and $0 \le r \le 1$, the third term must be positive for sufficiently large r, since $X_1^* - X_2^* < 1$. For any such r, the third term diverges as $\eta \to t$. Thus, $\frac{\partial \mathcal{W}_c}{\partial \sigma} > 0$ for sufficiently large η and r, in which case the default's cumulative welfare effect is positive. □

Proof of Proposition 14

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

$$\phi(\widetilde{z}_1^*) < \left(\frac{\mu_1'}{\mu_2'}\right)^2 \phi(\widetilde{z}_2^*), \tag{61}$$

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 (61). 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 (61) 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 (61) is true. □

Proof of Proposition 15

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

$$\begin{aligned} & u_1^{\delta_1\delta_2}(\widetilde{x}^{\delta_1\delta_2}) = u_2^{\delta_1\delta_2}(\widetilde{x}^{\delta_1\delta_2}) \\ \iff \quad & V_1 - \delta_2\sigma + X_1^{\delta_1\delta_2}\eta - h\alpha_1 - t\widetilde{x}^{\delta_1\delta_2} = V_2 - \delta_1\sigma + X_2^{\delta_1\delta_2}\eta - h\alpha_2 - t(1 - \widetilde{x}^{\delta_1\delta_2}), \end{aligned} \tag{62}$$

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

$$X_1^{\delta_1\delta_2} = \begin{cases} \widetilde{x}^{10} + w, & \text{if } \delta_1\delta_2 = 10 \\ \widetilde{x}^{01}, & \text{if } \delta_1\delta_2 = 01 \\ (1+w)\widetilde{x}^{00}, & \text{if } \delta_1\delta_2 = 00, \end{cases} \quad X_2^{\delta_1\delta_2} = \begin{cases} 1 - \widetilde{x}^{10}, & \text{if } \delta_1\delta_2 = 10 \\ 1 - \widetilde{x}^{01} + w, & \text{if } \delta_1\delta_2 = 01 \\ (1+w)(1-\widetilde{x}^{00}), & \text{if } \delta_1\delta_2 = 00. \end{cases} \tag{63}$$

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:

$$\widetilde{x}^{10} = \frac{1}{2} + \frac{V_1 - V_2 + \sigma + w\eta - h(\alpha_1 - \alpha_2)}{2(t-\eta)}, \quad \widetilde{x}^{01} = \frac{1}{2} + \frac{V_1 - V_2 - \sigma - w\eta - h(\alpha_1 - \alpha_2)}{2(t-\eta)}. \tag{64}$$

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

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

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 $p_i^* = \mu_i(\widetilde{z}_i^*)$ continues to apply in all cases in this extension. Additionally, when a default is in effect, the second FOC is also the same: $X_i^{\delta_1\delta_2 *} = \frac{h}{2(t-\eta)}\mu(\widetilde{z}_i^*)$. However, under a choice screen, the second FOC changes slightly: it becomes $\widetilde{x}^{00*} = \frac{h}{2(t-\eta)}\mu(\widetilde{z}_1^*)$ for platform 1 and $1 - \widetilde{x}^{00*} = \frac{h}{2(t-\eta)}\mu(\widetilde{z}_2^*)$ 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 $p_i^* = \mu(\widetilde{z}_i^*)$ thus implies $p_i^* = \widetilde{z}_i^*$ and thus $\alpha_i^* = p_i^* + \widetilde{z}_i^* = 2\widetilde{z}_i^*$. Plugging these expressions into the FOCs and solving the equations, we find the following equilibrium

results. First, when platform 1 is the default:

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

When platform 2 is the default:

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

Finally, under a choice screen:

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

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 16

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 3, 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.

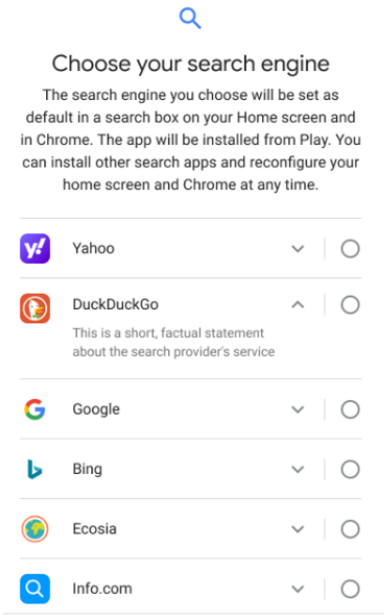


Figure B 1. The Android choice screen implemented in the EEA.

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:

$$\mathcal{Z}^* = \left\{(\widetilde{z}_1, \widetilde{z}_2) \,\middle|\, \mu(\widetilde{z}_1) + \mu(\widetilde{z}_2) = \tfrac{2(t-\eta)}{h},\ \mu(\widetilde{z}_1) \geq 0, \mu(\widetilde{z}_2) \geq 0\right\} \tag{66}$$

Using (11), (12), and (13), this locus contains all the different equilibria in which both platforms are active—one for each value of $\Delta \in (-\Delta^{crit}, \Delta^{crit})$.[58] 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

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

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 $(\widetilde{z}^{max}, \underline{z})$ when $\Delta = \Delta^{crit}$. (For simplicity, the figure assumes that $\underline{z} = 0$.)

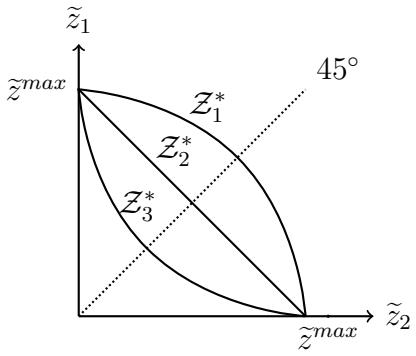


Figure B 2. Examples of $\mathcal{Z}^*$ corresponding to different functional forms of $\mu(\cdot)$.

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 $\widetilde{z}_1$ and a decrease in $\widetilde{z}_2$, but the relative magnitudes of these effects depend on the shape of μ. If μ is convex, then the change in $\widetilde{z}_1$ is small and the change in $\widetilde{z}_2$ 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 $\alpha_i = \widetilde{z}_i + p_i$, and we know from (13) that the changes in p_1 and p_2 are always equal in magnitude (that is, $\frac{\partial p_1}{\partial \Delta} = -\frac{\partial p_2}{\partial \Delta}$).

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.

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.[59] 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 Motorola. This raises the

[59]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.

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. (2025) 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.[60] 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 (for example, 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 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

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

consumer welfare. For an infinitesimal switching cost, this boils down to the requirement that

$$\frac{\partial \Pi_1^*}{\partial \sigma} \geq -\frac{\partial \mathcal{W}_c^*}{\partial \sigma} \tag{67}$$

Assuming $\eta \leq \frac{2}{3}t$, it is straightforward to verify that the following inequality is a necessary condition for (67):

$$\frac{m}{h}[1 + 2\mu_1'(\widetilde{z}_1^*)]F(\widetilde{z}_1^*)X_1^* \geq 1 + 3\mu'(\widetilde{z}_1^*)X_1^*. \tag{68}$$

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

B.4. Advertiser Participation and Consumer Utility

Models of two-sided platforms usually assume that each side's utility depends directly on the number of participating users on the other side. However, it is ad intensity (α_i), not advertiser participation (Z_i), that appears in the utility functions in (4). This is intended to better capture consumer preferences in the search market. As the Department of Justice noted in its *Google* complaint:

> As the number of users of a general search engine grows, advertisers benefit because they want their marketing campaigns to reach large groups of consumers. But users do not benefit from indirect network effects in an equivalent way. As Google's Chief Economist [Hal Varian] has explained, "users do not decide which search engine to use based on the number of advertisers."

To understand this, it is important to distinguish between (1) the total number of advertisers who use a given search platform; and (2) the number/intrusiveness of ads *on the search page*. It is mainly (2) that consumers care about. More generally, consumers care about the advertising activity they are directly exposed to while using a search engine (including gathering of their data or manipulation of search results). The ad intensity variable is designed to capture this.

Still, there is a sense in which consumer utility depends indirectly on advertiser participation (Z_i). Recall that the marginal advertiser type is $\widetilde{z}_i \equiv \alpha_i - p_i$. And, using (2), $Z_i = mF(\widetilde{z}_i)$. Hence, because F is invertible,[61] ad intensity can be decomposed as

$$\alpha_i = F^{-1}(Z_i/m) + p_i.$$

Thus, consumer disutility from advertising ($-h\alpha_i$) includes a contribution of $-hF^{-1}(Z_i/m)$. Accordingly, advertiser participation (Z_i) does affect consumer utility. However, utility does not depend on Z_i *independently* of α_i. Rather, Z_i affects utility through its influence on α_i.

[61] This follows from the assumption that F has interval support, which would not be true if F were constant over some region.