# Choice Screen Auctions

Michael Ostrovsky*
Stanford University and NBER
ostrovsky@stanford.edu

October 29, 2020

*PRELIMINARY AND INCOMPLETE, PLEASE DO NOT CIRCULATE*

**Abstract**

Choice screen auctions have recently been deployed in 31 European countries, allowing consumers to choose their preferred search engine on Google's Android platform instead of automatically defaulting them to Google's own search engine. I show that a seemingly minor detail in the design of these auctions—whether they are conducted on a "per appearance" or a "per install" basis—plays a major role in the mix and characteristics of auction winners, and consequently in their expected overall market share. I also show that "per install" auctions distort the incentives of alternative search engines toward extracting as much revenue as possible from each user who installs them, at the expense of lowering the expected number of such users. The distortion gets worse as the auction gets more competitive and the number of bidders increases. Empirical evidence from Android choice screen auctions conducted in 2020 is consistent with my theoretical results.

*The opinions expressed in this paper are my own, and do not necessarily reflect those of the organizations with which I am currently affiliated or have been affiliated in the past.

# 1 Introduction

Optimal regulation of digital platforms is one of the thorniest issues in competition policy. A particularly challenging dimension for regulation is the fact that dominant platforms are often active in multiple distinct businesses, and may leverage their position in one area into gaining an advantage in another. The net effect of such leverage on consumer welfare is often ambiguous and hard to determine. On one hand, a dominant platform's expertise and technological complementarities may make the adjacent product genuinely superior to the alternatives. On the other hand, such leverage may make it harder for other firms to successfully compete, even if their products, on their own, would be preferred by some consumers to that of the platform.

These linkages across product lines have led regulators to sometimes propose extreme measures to regulate large digital platforms, all the way to breaking them up and prohibiting them from entering certain lines of business. Notable examples in the U.S. include the Microsoft case of the late 1990s, in which the initial court decision was to order the company to "structurally reorganize,"[1] and the recently concluded congressional investigation into the business practices of Amazon, Apple, Facebook, and Google, which proposes "structural separations and line of business restrictions" as a solution for "restoring competition in the digital economy."[2] Regulators in the European Union and other parts of the world have often reached similar conclusions. Of course, a breakup of a company is a very heavy-handed solution, difficult to implement, rife with potential unintended consequences, and, unsurprisingly, adamantly opposed by the digital platforms.[3]

In light of these problems, in certain cases a more "lightweight" alternative has been adopted as a compromise solution by some of the platforms and regulators: choice screens. The logic of a choice screen is straightforward: instead of having a consumer use the dominant platform's product automatically and by default, the platform agrees to present to the consumer a menu of choices. This menu includes the platform's own product as one of the options, but also includes several competing products as alternatives. Consumers can then choose whichever products they prefer, leveling the playing field between the dominant platform and its competitors.

Choice screens for Web browsers on the Windows platform were first proposed by Microsoft in 1999 as a remedy in its negotiations with the U.S. Department of Justice.[4] They were not adopted

[1] `https://en.wikipedia.org/wiki/United_States_v._Microsoft_Corp`.

[2] "To address this underlying conflict of interest, Subcommittee staff recommends that Congress consider legislation that draws on two mainstay tools of the antimonopoly toolkit: structural separation and line of business restrictions. Structural separations prohibit a dominant intermediary from operating in markets that place the intermediary in competition with the firms dependent on its infrastructure. Line of business restrictions, meanwhile, generally limit the markets in which a dominant firm can engage." (October 6, 2020. See Section VI.A.I in `https://judiciary.house.gov/uploadedfiles/competition_in_digital_markets.pdf`.)

[3] To give just one recent example (out of many available ones): "A government effort to break up Facebook Inc. from Instagram and WhatsApp would defy established law, cost billions of dollars and harm consumers, according to a paper company lawyers have prepared in the wake of rising antitrust legal threats. [...] In the paper, Facebook says unwinding the deals would be nearly impossible to achieve, forcing the company to spend billions of dollars maintaining separate systems, weakening security and harming users' experience." (October 4, 2020, `https://www.wsj.com/articles/facebook-says-government-breakup-of-instagram-whatsapp-would-be-complete-nonstarter-11601803800`.)

[4] `https://www.wired.com/2000/11/microsoft-7/`

at that time, but were subsequently accepted as a compromise solution between the European Commission and Microsoft in 2009, and were displayed to users in Europe for five years, from 2010 until 2014.[5] (Due to a technical error, the choice screen was not displayed on one of the versions of Windows from May 2011 until July 2012, affecting approximately 15 million users. Microsoft admitted its responsibility for this error, and was subsequently fined €561 million for it.[6]) In 2017, Google reached a settlement with the competition authority in Russia to display choice screens for the default search engine on the Android platform there.[7] A similar agreement was reached between Google and the European Commission following a €4.3 billion fine imposed on the company by the Commisision in 2018,[8] and Google began displaying choice screens for both default search engines and web browsers to Android users in Europe in 2019.[9]

Choice screen menus can be an effective and powerful tool. For instance, following the 2010 introduction of browser choice screen menus on the Windows platform in Europe, the number of downloads of Opera Software's web browser more than doubled.[10] Discussing Google's introduction of choice screen menus on Android, the European Commissioner for Competition Margrethe Vestager stated, "We've seen in the past that a choice screen can be an effective way to promote user choice."[11] At the same time, from the point of a company that owns the platform, the initial implementations of choice screen menus suffered from one serious shortcoming: zero revenue. This

[5] "Under the commitments approved by the Commission, Microsoft will make available for five years in the European Economic Area [...] a "Choice Screen" enabling users of Windows XP, Windows Vista and Windows 7 to choose which web browser(s) they want to install in addition to, or instead of, Microsoft's browser Internet Explorer. ...

The Commission's preliminary view was that competition was distorted by Microsoft tying Internet Explorer to Windows. This was because it offered Microsoft an artificial distribution advantage not related to the merits of its product on more than 90 per cent of personal computers. Furthermore, the Commission's preliminary view was that this tying hindered innovation in the market and created artificial incentives for software developers and content providers to design their products or web sites primarily for Internet Explorer.

The approved commitments address these concerns. PC users, by means of the Choice Screen, will have an effective and unbiased choice between Internet Explorer and competing web browsers. This should ensure competition on the merits and allow consumers to benefit from technical developments and innovation both on the web browser market and on related markets, such as web-based applications." (December 16, 2009, `https://ec.europa.eu/commission/presscorner/detail/en/IP_09_1941`.)

[6] `https://ec.europa.eu/commission/presscorner/detail/en/IP_13_196`.

[7] `https://www.reuters.com/article/us-alphabet-google-russia-idUSKBN17J11C`, `https://yandex.com/blog/yacompany-com/choosing-yandex-search-on-android`.

[8] "The Commission decision has concluded that Google has engaged in two instances of illegal tying:

First, the tying of the Google Search app. As a result, Google has ensured that its Google Search app is pre-installed on practically all Android devices sold in the EEA. Search apps represent an important entry point for search queries on mobile devices. The Commission has found this tying conduct to be illegal as of 2011, which is the date Google became dominant in the market for app stores for the Android mobile operating system.

Second, the tying of the Google Chrome browser. As a result, Google has ensured that its mobile browser is pre-installed on practically all Android devices sold in the EEA. Browsers also represent an important entry point for search queries on mobile devices and Google Search is the default search engine on Google Chrome. The Commission found this tying conduct to be illegal as of 2012, which is the date from which Google has included the Chrome browser in its app bundle."

(July 18, 2018. `https://ec.europa.eu/commission/presscorner/detail/en/IP_18_4581`.)

[9] `https://www.blog.google/around-the-globe/google-europe/presenting-search-app-and-browser-options-android-users-europe/`.

[10] `https://press.opera.com/2010/03/18/opera-more-than-doubles-download-numbers-in-europe-after-choice-screen-introduction/`.

[11] `https://ec.europa.eu/commission/presscorner/detail/en/STATEMENT_19_1774`.

may be a particularly salient issue in the case of search engines, where being chosen by a consumer is extremely valuable to a search engine due to the advertising revenues it expects to receive when the consumer uses it, and where, moreover, the dominant company itself may have to make large payments to *another* platform to have consumers use its search engine there.[12] In this case, it is logical for the company to argue that it should be allowed to charge others for the right to have their products be shown on its platform's choice screens—and a natural way to do so is via an auction.

That is the decision that Google announced in August 2019,[13] and the first "choice screen auctions" took place in early 2020. The basic rules of Google's choice screen auctions are very simple.

> In each country auction, search providers will state the price that they are willing to pay each time a user selects them from the choice screen in the given country. The three highest bidders will appear in the choice screen for that country. The provider that is selected by the user will pay the amount of the fourth-highest bid.[14]

In the same document, Google explains why it chose to auction off slots in the choice screen this way:

> Q: Why does Google use an auction to determine the search providers that appear in the choice screen?
>
> A: An auction is a fair and objective method to determine which search providers are included in the choice screen. It allows search providers to decide what value they place on appearing in the choice screen and to bid accordingly.
>
> The auction revenues help us to continue to invest in developing and maintaining the Android platform.

In this paper, I show that a seemingly minor detail of the implementation of choice screen auctions plays a major role in their outcomes—and thus in the overall effectiveness of the antitrust remedy. Specifically, while the answer in the Q&A section of the document says that an auction "allows search providers to decide what value they place on appearing in the choice screen and to bid accordingly," the auction, as implemented, charges these providers not for *appearing* in the choice screen but for *being chosen by a user*.

While the difference may seems to be just a matter of language, it is not. To see the intuition for the difference, consider a version of the auction with just one available spot and two bidders. Bidder A gets revenue \$10 from each user who installs its search engine, and if it is shown as an option in the choice screen, the probability that a user will choose it is 10%. Bidder B gets revenue

[12] While the companies do not directly disclose these numbers, analysts estimate that Google is paying Apple an amount on the order of \$10 billion dollars per year to have Google be the default search engine on Apple's Safari browser (`https://finance.yahoo.com/news/google-pay-apple-12-billion-155007747.html`).

[13] `https://www.blog.google/around-the-globe/google-europe/update-android-search-providers-europe/`

[14] `https://www.android.com/choicescreen/`

$20 from each user who installs its search engine, but the probability that a user will choose it (if it is shown as an option in the choice screen) is only 1%. The value bidder A has for appearing on the screen is therefore $1, and the value that bidder B has for appearing on the screen is $0.20. Thus, if the auction was conducted on the "per appearance" basis, bidder A would win, would pay $0.20 per appearance, and would have its search engine be chosen by users 10% of the time, while the dominant platform's own search engine would be chosen 90% of time time. If, instead, the auction is conducted as implemented, with bidding and payment on the "per install" basis, bidder B will win and will pay $10 every time its search engine is chosen—or $0.10 per appearance. The winner's search engine will be chosen only 1% of the time, and the dominant platform's one will be chosen the remaining 99% of time. Thus, relative to the per-appearance auction, the per-install auction results in a lower likelihood that an alternative search engine will be chosen by the user (making it correspondingly more attractive to the dominant platform) and gives advantage to search engines that generate higher revenue per user vs. those that are more popular but generate less revenue on a per-user basis. In Section 2, I show that these conclusions hold more generally, in a basic model in which alternative search engines differ on two dimensions: revenue per user (i.e., how much revenue the search engine generates, on average, when a user chooses to install it) and popularity (i.e., the likelihood that the search engine will be chosen if it is shown to the user in the choice screen). Moreover, I show that the difference gets exacerbated by competition. As the number of alternative search engines grows, under the per-appearance auction, the expected popularity of the winner also grows, and the probability that the dominant platform's own search engine is chosen decreases. By contrast, under the per-install auction, these measures are not affected by the number of bidders.

In the example above and in the model of Section 2, a search engine's popularity and revenue-per-user (RPU) are fixed. In practice, a search engine has some ability to trade them off against each other. For instance, a search engine may choose to show more intrusive ads, increasing its revenue per user but decreasing its popularity. Conversely, a search engine may donate some of its proceeds to charity or implement very strict privacy rules, lowering its revenue per user but increasing the probability that a user will choose it. I introduce this possibility in Section 3, and show that the two auction formats result in very different incentives to the search engines regarding this tradeoff. Under the per-appearance auction, each bidder chooses the same point on the popularity–RPU frontier as it would choose if it was the only bidder (and was thus guaranteed the spot on the choice screen). In particular, this implies that just as in the model of Section 2, as the number of competitors increases, the expected popularity of the winner also increases, and the expected probability that the dominant platform's search engine is chosen goes down. By contrast, under the per-install auction, each bidder has a strong incentive to distort the choice toward higher revenue per each user who chooses the product, at the expense of lowering the probability of actually being chosen. This distortion grows stronger as the number of bidders grows, and in the limit as that number approaches infinity, the distortion results in a "race to the bottom," with all bidders pushing to the extreme point on the popularity–RPU frontier: the highest possible RPU and the lowest possible popularity. As a result, the expected popularity of the winner goes in the opposite

direction vs. that in the case of the per-appearance auction, minimizing the probability that an alternative search engine will be chosen.

In Section 4, I present empirical evidence from the first three sets of choice screen auctions (for the periods of March–June 2020, July–September 2020, and October–December 2020) conducted by Google in 31 European countries. The evidence is consistent with my theoretical conclusions. In particular, to mention just one data point from the section, the search engine that was most successful in these auctions, winning a slot in every country and in every period, has been installed less than 50,000 times *worldwide* and has so few user ratings that Google's Play Store does not provide statistics on those ratings. For the sake of comparison, Google's own search app has been installed more than 5 billion times, while alternative search engines Bing (produced by Microsoft) and DuckDuckGo (produced independently and focused on user privacy) have been installed more than 10 million times. (These numbers include all installs, including those that come from choice screens and those that do not.)

Section 5 concludes.

## 2 Basic Model: Exogenous Popularity and Revenue-per-User

A platform is auctioning off the right to be shown on the choice screen. There is only one slot available, next to the platform's own product.[15] There are $n = 2$ bidders, $i \in 1, 2$. Each bidder $i$ has an exogenously determined *popularity* $q_i$ and if its product is chosen by a user, bidder $i$ receives revenue $r_i$ from that. Variables $q_i$ and $r_i$ are private information of bidder $i$. Variables $q_1$, $q_2$, $r_1$, and $r_2$ are independently and identically distributed, and each is drawn from the uniform distribution on $[0, 1]$.

If a product of popularity $q$ is shown to a user, it is chosen with probability $q$. The platform gets benefit $\pi > 1$ if its own product is chosen, and the auction payment if the user chooses an alternative.[16]

Under the "per appearance" auction, each bidder submits a bid for the right to be shown to users. The bidder with the highest bid wins and is shown on the choice screen, and pays the amount equal to the bid of the second-highest bidder.

Under the "per install" auction, each bidder submits a bid. The bidder with the highest bid wins and is shown on the choice screen, and pays the amount equal to the bid of the second-highest bidder *if the user chooses its product.*

Note that both the "per appearance" and "per install" auctions are incentive compatible: it is a dominant strategy for each bidder to submit its valuation truthfully.[17] That is, under the "per install" auction, each bidder $i$ will bid $r_i$, while under the "per appearance" auction each bidder will bid $q_i r_i$.

[15] In Google's Android choice screen auctions, there are three slots next to the advertiser's own listing. I consider the case of only one slot for simplicity; this assumption does not qualitatively change my conclusions.

[16] The idea behind this assumption is that a platform's first choice is to have its own product chosen—after all, if that was not the case, there would be no need for the antitrust remedy.

[17] I will ignore other equilibria of these auctions.

To characterize the distribution of outcomes in the per-appearance auction, we need to perform some calculations.

First, observe that the unconditional distribution of each bid $b_i = q_i r_i$ is given by the CDF $G(x) = x - x\ln x$ and the corresponding density function $g(x) = -\ln x$ (for $x \in (0,1]$).[18] Thus, for a bidder with type $(q,r)$, the probability of winning the auction is $(qr) - (qr)\ln(qr)$. For the population of bidders with type $q$ and with types $r$ distributed uniformly on $[0,1]$, the probability of winning is therefore $\int_0^1 ((qr) - (qr)\ln(qr))\,dr = q\int_0^1 (r - r\ln q - r\ln r)\,dr = q\left(\frac{1}{2} - \frac{1}{2}\ln q + \frac{1}{4}\right) = \frac{3}{4}q - \frac{1}{2}q\ln q$. We can now calculate the expected popularity of the winner of the auction, which is equal to $2\int_0^1 q\left(\frac{3}{4}q - \frac{1}{2}q\ln q\right) = \frac{11}{18}$.[19] Note that since $q$ and $r$ play symmetric roles in determining the winner and have the same distribution, the expected revenue per install of the winner will also be equal to $\frac{11}{18}$.

The expected payment made by the winner of the auction is equal to $E[\min\{q_1 r_1, q_2 r_2\}]$. Given the distribution $G(\cdot)$ of each $q_i r_i$ derived above, the distribution of $\min\{q_1 r_1, q_2 r_2\}$ is given by $G(x)^2 + 2G(x)(1-G(x))$, with the corresponding density $2g(x) - 2g(x)G(x) = -2\ln x(1 - x + x\ln x)$. Thus, $E[\min\{q_1 r_1, q_2 r_2\}] = -2\int_0^1 (x\ln x(1 - x + x\ln x))\,dx = \frac{7}{54}$, and the expected payoff of the platform is $\frac{7}{54} + \frac{7}{18}\pi$.

Characterizing the outcomes of the per-install auction is straightforward. Under this format, the expected popularity of the winner is $1/2$—it is independent of the bids in the auction and plays no role in determining the winner. This is lower than the expected popularity of the winner in the per-appearance auction ($\frac{11}{18}$), and thus the probability that an alternative product will be chosen is reduced and the probability that the dominant platform's product is chosen is increased.

The expected per-install payment made by the winner is $1/3$ (and is independent of its popularity), and the expected payoff of the platform in the per-install auction is therefore $1/6 + 1/2\pi$, which is higher than its expected payoff in the per-appearance auction ($\frac{7}{54} + \frac{7}{18}\pi$). The expected revenue per install of the winner is $2/3$, which is also higher than the corresponding quantity in the per-appearance auction ($\frac{11}{18}$).

The contrast between the outcomes of the per-appearance and per-install auctions becomes even more striking if the number of bidders, $n$, becomes large—i.e., the number of potential alternative products grows. It is immediate that as $n$ grows, the expected popularity of the winner of the per-appearance auction converges to one—the highest possible popularity. By contrast, the expected popularity of the winner of the per-install auction remains unchanged, at one half—under that format, the increase in the competition has no impact on the popularity of the winner, and thus on the probability that an alternative product will be chosen. Of course, if the platform's payoff from having a user choose its own product is higher than the revenue that an alternative product generates, this outcome is preferred by the platform, just as it was in the case of $n = 2$.

[18] For $x \in [0,1]$, the probability that $q_i r_i$ is less than or equal to $x$ is equal to $G(x) = x + \int_x^1 \frac{x}{q}dq = x + x(\ln 1 - \ln x) = x - x\ln x$.

[19] The logic behind this formula is that by symmetry, (the expected popularity of the winner of the auction) is equal to (the expected popularity of the winner of the auction conditional on that winner being bidder 1). The latter expression is equal to $\frac{\int_0^1 q_1 Prob(\text{bidder 1 is the winner of the auction}|q_1)dq_1}{Prob(\text{bidder 1 is the winner of the auction})}$, whose denominator is equal to $\frac{1}{2}$.

# 3 Extension: Endogenous Popularity and Revenue-per-User

It is clear from the results of Section 2 that the choice between per-install and per-appearance auction formats is a first-order issue. However, that only tells a part of the story. In this section, I consider the incentives of bidders to choose between making their product more attractive to the users (at the expense of lower per-user revenue) vs. moving in the opposite direction. This is a an important issue for the case of default search engines on a platform. Search engines may be able to increase their popularity by reducing the intrusiveness of ads, enhancing privacy protections, or donating their advertising revenue to charitable causes. These measures, while making the search engine more attractive to the users, reduce the revenue it makes from each one of them. In this section I show that the issues discussed in Section 2 are exacerbated once these incentives are taken into account—in some cases, dramatically so.

I maintain most of the assumptions of Section 2, and make only one change. Instead of assuming that each bidder $i$ has an exogenously given popularity $q_i$ and revenue-per-user $r_i$, I assume that each bidder has an exogenously given type $t_i$ drawn from independently from the uniform distribution $F(t)$ on $[0,1]$ (with the corresponding density function $f(t)$). The bidder can then select its popularity $q_i$ from $[0,t_i]$, and its revenue-per-user is then equal to $r_i = t_i - q_i$. After making this decision, the bidder submits its auction bid, as before.

Bidding decisions are still straightforward under both rules—it is optimal to bid truthfully in both per-appearance and per-install auctions. However, each bidder now needs to decide, given its type $t_i$, how much of that type to allocate to popularity $q_i$, and how much to allocate to revenue-per-user. I work out the resulting equilibria in the next two subsections.

## 3.1 Per-appearance Auction

Consider a per-appearance auction with $n$ bidders, and suppose bidder $i$ has type $t_i \in [0,1]$. Fix other bidders' strategies, let $G(x)$ denote the distribution of the first-order statistic of those bidder's bids, and let $P(x)$ denote the expected payment that bidder $i$ would make, conditional on winning the auction, if it submitted bid $x$ (note that $G(x)$ and $P(x)$ are purely functions of the other bidders' strategies and $x$). Bidder $i$ has two decisions to make: popularity $q_i$ and bid $b_i$. Its payoff, as a function of these two decisions, is given by

$$\Pi(q_i, b_i) = G(b_i) \times \left(q_i(t_i - q_i) - P(b_i)\right).$$

It is immediate that bidder $i$'s optimal choice of popularity is to set

$$q_i = \frac{t_i}{2}. \tag{1}$$

This is the optimal strategy *regardless* of what other bidders' strategies are (or how many of those bidders there are), and thus the strategy profile in which each bidder sets $q_i = t_i/2$ and then bids $(t_i/2)^2$ per appearance constitutes an equilibrium.

### 3.2 Per-install Auction

Equilibrium characterization in the case of the per-install auction requires a somewhat more subtle argument. Consider a symmetric equilibrium of the per-install auction with $n$ bidders, and suppose equilibrium strategies are given by functions $q(t)$ and $b(t)$, with the first one denoting the popularity chosen by a bidder with type $t$ and the second one denoting its bid. We know that in equilibrium, function $b(t)$ will be truthful (given the choice of popularity $q(t)$ and the corresponding revenue-per-user $t - q(t)$), but, just as in the case of the per-appearance auction, it is more convenient to not yet impose that restriction on function $b(t)$.

Take a bidder of type $t_i \in [0,1]$ and a real number $\Delta_q$ such that $t_i + \Delta_q \in [0,1]$. Let $\Pi(\Delta_q; t_i)$ denote the expected payoff of bidder $i$ whose type is $t_i$ *if it chooses popularity* $q(t_i + \Delta_q)$ *but bids* $b(t_i)$, given that other bidders are bidding according to strategies $q(t)$ and $b(t)$. We then have

$$\Pi(\Delta_q; t_i) = F^{n-1}(t_i) \times q(t_i + \Delta_q) \times \left( t_i - q(t_i + \Delta_q) - E\left[ b(\max_{j \neq i}\{t_j\} | \max_{j \neq i}\{t_j\} \leq t_i) \right] \right).$$

Since we started with an equilibrium profile of strategies, the partial derivative of $\Pi(\Delta_q; t_i)$ with respect to $\Delta_q$ has to be equal to zero when evaluated at $(0; t_i)$. This implies the following equation for $q(t_i)$:

$$q(t_i) = \frac{t_i - E\left[b(\max_{j \neq i}\{t_j\}) | \max_{j \neq i}\{t_j\} \leq t_i\right]}{2}. \tag{2}$$

Note that even without fully characterizing the equilibrium, from the comparison of equations (1) and (2), it is immediate that equilibrium popularity chosen by each type $t_i > 0$ will be strictly lower under the per-install auction than under the per-appearance auction (assuming, of course, that the equilibrium of the per-install auction actually exists, which we will show below).

Recall that by incentive compatibility, we have $b(t_j) = t_j - q(t_j)$. We can then rewrite the expectation in equation (2) as

$$\begin{aligned} E\left[ b(\max_{j \neq i}\{t_j\}) | \max_{j \neq i}\{t_j\} \leq t_i \right] &= \frac{\int_0^{t_i} (s - q(s)) dF^{n-1}(s)}{F^{n-1}(t_i)} \\ &= \frac{(n-1)\int_0^{t_i} (s - q(s)) f(s) F^{n-2}(s) ds}{F^{n-1}(t_i)}, \end{aligned}$$

and subsequently rewrite equation (2) as

$$2q(t_i)F^{n-1}(t_i) = t_i F^{n-1}(t_i) - (n-1)\int_0^{t_i} (s - q(s)) f(s) F^{n-2}(s) ds. \tag{3}$$

The next step is to take a derivative of both sides of equation (3) with respect to $t_i$, which gives us

$$\begin{aligned} 2q'(t_i)F^{n-1}(t_i) + 2(n-1)q(t_i)f(t_i)F^{n-2}(t_i) &= F^{n-1}(t_i) + (n-1)t_i f(t_i) F^{n-2}(t_i) \\ &- (n-1)t_i f(t_i) F^{n-2}(t_i) + (n-1)q(t_i) f(t_i) F^{n-2}(t_i), \end{aligned}$$

which simplifies to

$$2q'(t_i) + (n-1)q(t_i)\frac{f(t_i)}{F(t_i)} = 1. \tag{4}$$

Equation (4) is a linear differential equation, with the initial condition $q(0) = 0$. In our case, $F(t_i) = t_i$ and $f(t_i) = 1$, and so the equation becomes

$$2q'(t_i) + (n-1)q(t_i)\frac{1}{t_i} = 1,$$

with the solution

$$q(t_i) = \frac{t_i}{n+1}.$$

### 3.3 Comparison

The comparison between the two formats is immediate. Even with just two bidders, $n = 2$, a bidder of each type $t_i$ chooses a much lower popularity under the per-install auction than under the per-appearance auction: $\frac{t_i}{3}$ vs. $\frac{t_i}{2}$. Of course, this lower popularity on a per-type basis immediately translates into a correspondingly lower overall probability that a product alternative to the dominant platform's one will be picked by users from the choice screen.

With more bidders, the difference becomes even more dramatic. Under the per-appearance auction format, a bidder's choice of the point on the popularity–RPU frontier is unaffected by the number of bidders, and thus as $n \to \infty$, the popularity of the winning bidder converges to $\frac{1}{2}$. By contrast, under the per-install format, the more bidders the auction has, the lower is the popularity that each of them chooses in equilibrium, and so as $n \to \infty$, the popularity of the winner converges to zero, and thus the share of installs that goes to the dominant platform's own product converges to 100%, completely undoing choice screen auctions' raison d'être.

## 4 Evidence from Android Choice Screen Auctions

I now turn to the empirical evidence on the outcomes of Google's choice screen auctions for default search engines on the Android platform. These auctions were conducted in January 2020 (for the period from March until June), June 2020 (for the period from July until September), and September 2020 (for the period from October until December). For each time period, 31 independent auctions were conducted (one per country). These auctions were conducted on a per-install basis, with those submitting top three bids being shown on the choice screen and paying the fourth highest price every time a user chose one of them from the choice screen. In the event of a tie, more than three bidders could win, with the ties broken randomly on a per-device basis.

Google lists latest auction winners at `https://www.android.com/choicescreen-winners/`, and the lists of winners from the earlier auctions are available on the corresponding webpage on the Internet Archive.[20] These results, on a country-by-country basis, are summarized in Table 1.

[20] `https://web.archive.org/web/*/https://www.android.com/choicescreen-winners/`

In the table, countries are sorted by population (from top to bottom), while search engines are sorted by the average population they won across the three time periods. E.g., the search engine DuckDuckGo won a slot in every country in the first period (total population 519.4 million), and also in the second period (same total population), but only four countries in period 3 (Bulgaria, Croatia, Iceland, and Liechtenstein, with the total population 11.4 million). The average population won by DuckDuckGo across the three time periods is thus $(519.4 + 519.4 + 11.4)/3 = 350.1$ million, as listed in the last row of Table 1. Finally, in each search engine–country cell I list the periods in which that search engine won a slot on the choice screen menu in that country.

Table 2 lists the search engines by the average population they won across the three periods, and adds data collected from the Android Play Store (`https://play.google.com/store/apps`) on the popularity and quality of these search engines.[21] To be eligible to participate in a choice screen auction, a "search provider must have an app that is available for free in Google Play," and if a search engine is chosen by a user from the choice screen, "[that app] will be downloaded from Play," in addition to the search engine being set as the default in the Chrome browser on the user's device (`https://www.android.com/choicescreen/`). Of course, users can install these apps even without the choice screen (as they do, e.g., outside of Europe), and moreover, the numbers listed on Play Store count the number of installs worldwide, not just in Europe. Thus, while the numbers of downloads resulting from choice screen auctions are not publicly disclosed by Google or the European Commission, the install numbers from Play Store, listed in Table 2, provide upper bounds on those numbers.[22] These install numbers, of course, also provide a measure of the overall popularity of these search engines, by including the installs that were made independently of choice screen auctions.[23]

Table 2 also includes the ratings from Android users on the quality of these apps. Play Store requires a minimum number of installs, ratings, and reviews before it starts disclosing information about average ratings to the users. Three of the ten winners of choice screen auctions did not pass this bar, and thus do not have average ratings reported by the Play Store platform (indicated as "n/a" in the table).

Consistent with the theoretical results from Section 2, there is no meaningful correlation between the quality and popularity of the search engines participating in these auctions and the average population they won in them. Strikingly, of the top 5 winners, each of whom on average won a slot in choice auctions covering a population of more than 100 million, three had fewer than 50,000 installs (from all sources, over their lifetimes), and two, including the search engine that won the most coverage in the auction, had so few user ratings that Google did not provide their averages.[24]

[21] I collected the data on September 30, 2020, the last day of the second choice screen auction period. The screenshots of Play Store pages from which these data were taken are available at `https://web.stanford.edu/~ost/papers/screenshots.html`.

[22] Note that Play Store does not disclose the exact number of installs, but instead shows a range, which is what Table 2 reports.

[23] One caveat on this measure is that the search engines' popularity may vary by geography. For example, the majority of users of the Yandex search engine are from Russia and the countries of the former Soviet Union.

[24] info.com had been rated a total of 18 times, and GMX had been rated 9 times. By comparison, DuckDuckGo had been rated 637,387 times.

Table 1: Winners of Android Choice Screen Auctions (by country)

| Country | Population | info.com | DuckDuckGo | PrivacyWall | Bing | GMX | Qwant | Yandex | Seznam | Givero | Ecosia |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Germany | 83,990,646 | 1, 2, 3 | 1, 2 | 2, 3 | 3 | 1 | | | | | |
| United Kingdom | 67,999,326 | 1, 2, 3 | 1, 2 | 2, 3 | 1, 3 | | | | | | |
| France | 65,297,182 | 1, 2, 3 | 1, 2 | 2, 3 | 3 | | 1 | | | | |
| Italy | 60,457,546 | 1, 2, 3 | 1, 2 | 2, 3 | 3 | | 1 | | | | |
| Spain | 46,789,532 | 1, 2, 3 | 1, 2 | 2, 3 | 3 | | 1 | | | | |
| Poland | 37,847,219 | 1, 2, 3 | 1, 2 | 3 | | 2, 3 | | 1 | | | |
| Romania | 19,196,044 | 1, 2, 3 | 1, 2 | 1, 3 | | 2, 3 | | 1, 3 | | | |
| Netherlands | 17,142,323 | 1, 2, 3 | 1, 2 | 2, 3 | 3 | 1 | | | | | |
| Belgium | 11,608,284 | 1, 2, 3 | 1, 2 | 2, 3 | 3 | | 1 | | | | |
| Czech Republic | 10,717,516 | 1, 2, 3 | 1, 2 | 3 | | | | | 1, 2, 3 | | |
| Greece | 10,415,204 | 1, 2, 3 | 1, 2 | 1, 3 | | 2, 3 | 1 | 1, 3 | | | |
| Portugal | 10,190,296 | 1, 2, 3 | 1, 2 | 1 | | 2, 3 | 1 | 1, 3 | | | |
| Sweden | 10,114,623 | 1, 2, 3 | 1, 2 | 1, 2, 3 | 3 | | | | | | |
| Hungary | 9,657,366 | 1, 2, 3 | 1, 2 | 1, 3 | | 2, 3 | | 1, 3 | | | |
| Austria | 9,032,162 | 1, 2, 3 | 1, 2 | 2, 3 | 3 | 1 | | | | | |
| Bulgaria | 6,938,828 | 1, 2, 3 | 1, 2, 3 | 1, 3 | | 2, 3 | | 1, 3 | | | |
| Denmark | 5,795,666 | 1, 2, 3 | 1, 2 | 1, 2, 3 | 3 | | 1 | 1 | | 1 | |
| Finland | 5,543,674 | 1, 2, 3 | 1, 2 | 2, 3 | 3 | | | 1 | | | |
| Slovakia | 5,461,816 | 1, 2, 3 | 1, 2 | | | 3 | | | 1, 2, 3 | | |
| Norway | 5,428,345 | 1, 2, 3 | 1, 2 | 1, 2, 3 | 3 | | | | | | |
| Ireland | 4,960,177 | 1, 2, 3 | 1, 2 | 1, 2, 3 | 3 | | | | | | |
| Croatia | 4,099,199 | 1, 2, 3 | 1, 2, 3 | 1 | | 2, 3 | | 1 | | | |
| Lithuania | 2,710,479 | 1, 2, 3 | 1, 2 | 3 | | 2, 3 | | 1 | | | |
| Slovenia | 2,079,635 | 1, 2, 3 | 1, 2 | 1 | | 2, 3 | | 1 | | | 3 |
| Latvia | 1,881,006 | 1, 2, 3 | 1, 2 | 3 | | 2 | | 1, 3 | | | |
| Estonia | 1,328,929 | 1, 2, 3 | 1, 2 | | | 2, 3 | | 1, 3 | | | |
| Republic of Cyprus | 1,190,962 | 1, 2, 3 | 1, 2 | 1 | | 2, 3 | | 1, 3 | | | |
| Luxembourg | 629,798 | 1, 2, 3 | 1, 2 | | | 3 | 1, 2, 3 | | | | |
| Malta | 514,564 | 1, 2, 3 | 1, 2 | 1, 3 | | 2, 3 | | 1 | | | |
| Iceland | 341,834 | 1, 2, 3 | 1, 2, 3 | 1 | | 2, 3 | | 1 | | | |
| Liechtenstein | 38,150 | 1, 2, 3 | 1, 2, 3 | | | 2, 3 | 1 | | | | |
| Average pop. (M) | | 519.4 | 350.1 | 326.4 | 154.1 | 110.4 | 70.8 | 60.2 | 16.2 | 1.9 | 0.7 |

Table 2: Popularity and Ratings of Android Search Engine Apps

| Search Engine | Av. Pop. | # Installs | Rating |
|---|---|---|---|
| info.com | 519.4 | 10,000 – 49,999 | n/a |
| DuckDuckGo | 350.1 | 10,000,000 – 49,999,999 | 4.7 |
| PrivacyWall | 326.4 | 10,000 – 49,999 | 3.9 |
| Bing | 154.1 | 10,000,000 – 49,999,999 | 4.5 |
| GMX | 110.4 | 5,000 – 9,999 | n/a |
| Qwant | 70.8 | 1,000,000 – 4,999,999 | 3.9 |
| Yandex | 60.2 | 100,000,000 – 499,999,999 | 4.5 |
| Seznam | 16.2 | 1,000,000 – 4,999,999 | 4.1 |
| Givero | 1.9 | 100 – 499 | n/a |
| Ecosia | 0.7 | 5,000,000 – 9,999,999 | 4.6 |

With regards to the theoretical results of Section 3 on endogenous quality choice, I do not have direct evidence on the impact of the incentives from the auction on the actual design choices that search engines make to trade off revenue they make per each user vs. their popularity and quality. The time period of less than a year may also be insufficient for these design changes to manifest themselves. However, these incentives are clearly understood (and lamented) by at least some of the participants in these auctions.

The most highly rated search engine, DuckDuckGo, won a slot in every country in periods 1 and 2, but then by period 3 only won slots in four countries covering in total only 2% of the available population. As explained by DuckDuckGo, "Despite DuckDuckGo being robustly profitable since 2014, we have been priced out of this auction because we choose to not maximize our profits by exploiting our users. In practical terms, this means our commitment to privacy and a cleaner search experience translates into less money per search. This means we must bid less relative to other, profit-maximizing companies."[25]

The second most highly rated search engine is Ecosia, which uses its profits to plant trees around the world (`https://info.ecosia.org/`). Despite its overall popularity with users (more than 5,000,000 installs), it only won coverage in the last period and only one country (covering less than 0.4% of overall available population). Ecosia explained its decision not to participate in the initial auction (for the first time period) as follows: "Ecosia is a not-for-profit search engine. Taking part in Google's auction would force us to spend our income on an unnecessary bidding war with other (profit-oriented) search engines. We'd rather use it to plant trees on our endangered planet."[26]

[25] `https://spreadprivacy.com/search-preference-menu-duckduckgo-elimination/`.
[26] `https://blog.ecosia.org/google-auction-choice-screen/`.

# 5 Discussion

It may be tempting to use the evidence like that in Section 4 to argue that choice screen auctions should be scrapped, and replaced with an alternative remedy. My theoretical results, however, show that such a conclusion may be premature. The issues described in Section 4 arise not from the usage of auctions per se, but rather from the specific implementation of these auctions. Replacing the "per install" bidding rule in choice screen auctions with a "per appearance" one may be sufficient to alleviate many of the shortcomings of the current implementation, or at least meaningfully reduce their impact.

A "per appearance" auction can be implemented in practice in a number of different ways. Search engines can bid directly per appearance, specifying how much they are willing to pay every time they are shown on the choice screen, and paying the highest losing bid every time they are shown. To help them make bidding decisions, the platform could provide them historical data on "conversion probabilities," i.e., the likelihood that the search engine will be chosen by a user conditional on being shown on the choice screen. Alternatively, the auction could still ask bidders to report their values "per install," but then rank bidders by the product of those reported values and estimated conversion probabilities and charge each winning bidder the lowest bid they could have reported and still won the auction. This approach is widely used in sponsored search auctions, and is well understood. It is straightforward for the platform to estimate these conversion probabilities. It can allocate a small fraction of "choice screen" impressions to display random selections of all eligible search engines, and use data from this fraction of traffic to estimate conversion probabilities.[27]

Choice screen auctions can be a powerful tool for leveling the playing field in the now widespread situations in which a dominant platform offers a product that is competing with several others. My results show that seemingly minor aspects of their implementation may have profound effects on their outcomes. More generally, the tools of market design can be useful for constructing elegant and effective regulatory solutions for the increasingly complex and interconnected digital economy. When deploying such solutions, it is important to analyze the equilibrium properties of resulting systems, on a very detailed level. Details matter.

[27] Google has a number of requirements for search engines to be eligible to participate in the auction, and any search engine that passes these requirements is deemed to be of sufficient quality to be offered as a choice to the users. E.g., in the official rules, Google says, "In the event that fewer than three eligible search providers bid, Google will fill any remaining slots randomly from the pool of eligible search providers on a per device basis. The pool of eligible providers will include those that applied to participate in the choice screen but did not submit bids." (`https://www.android.com/choicescreen/`)