NBER WORKING PAPER SERIES

SOURCES OF MARKET POWER IN WEB SEARCH: EVIDENCE FROM A FIELD EXPERIMENT

Hunt Allcott
Juan Camilo Castillo
Matthew Gentzkow
Leon Musolff
Tobias Salz

Working Paper 33410

http://www.nber.org/papers/w33410

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2025

We are grateful to Marc Cartright, Anlei Dong, Glenn Ellison, Chiara Farronato, Jingyi Guo, Paul Heidhues, Minha Hwang, Aadharsh Kannan, Widad Machmouchi, Markus Mobius, Sarah Moshary, Aviv Nevo, Michael Schwartz, Fiona Scott Morton, Steve Tadelis, Catherine Tucker, Mike Whinston, and audiences at Charles River Associates, Chicago Booth, NBER SI Digital Economics, Harvard-MIT IO Seminar, HBS Markets and Competition Conference, Northwestern Antitrust Conference, NYU Stern IO day, Stanford SITE, Stanford Behavioral Seminar, University of Michigan, Wharton, UVA, VIDE Seminar, Temple University, and TSE Economics of Platforms Seminar for helpful comments. We thank Chris Karr and Audacious Software for dedicated work on the Search Extension browser extension, and we thank Chiara Farronato and Andrey Fradkin for allowing Search Extension to use code developed for their work. We thank Shotaro Beppu, Jack Cenatempo, Juan Carlos Cisneros, Grace Coogan, Sameer Nair-Desai, Shiqi Yang, Wanxi Zhou and especially Wonjoon Choi for exceptional research assistance. We are grateful to the Sloan Foundation, the Sloan Research Fellowship, the Stanford Institute for Economic Policy Research (SIEPR), and the Business¸‚ Economics and Public Policy Department at the Wharton School of the University of Pennsylvania for generous support. We also gratefully acknowledge Microsoft Bing for sharing data with us. The experiment was approved by the MIT Committee on the Use of Human Subjects (Protocol # 2308001088) and was registered in the American Economic Association Registry for randomized control trials under trial AEARCTR-0012884; the pre-analysis plan is available from https://www.socialscienceregistry.org/trials/12884. Disclosures: Allcott, Castillo, and Musolff have previously worked at Microsoft Research. Gentzkow has done litigation consulting for Google and has been a member of the Toulouse Network for Information Technology, a research group funded by Microsoft. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w33410

Sources of Market Power in Web Search: Evidence from a Field Experiment
Hunt Allcott, Juan Camilo Castillo, Matthew Gentzkow, Leon Musolff, and Tobias Salz
NBER Working Paper No. 33410
January 2025
JEL No. L4, L86

ABSTRACT

We evaluate the economic forces that contribute to Google's large market share in web search. We develop a model of search engine demand in which consumer choices are influenced by switching costs, quality beliefs, and inattention, and estimate it using a field experiment with US desktop internet users. We find that (i) requiring Google users to make an active choice among search engines increases Bing's market share by only 1.1 percentage points, implying that switching costs play a limited role; (ii) Google users who accept our payment to try Bing for two weeks update positively about its relative quality, with 33 percent preferring to continue using it; and (iii) after changing the default from Google to Bing, many users do not switch back, consistent with persistent inattention. In our model, correcting beliefs and removing choice frictions would increase Bing's market share by 15 percentage points and increase consumer surplus by $6 per consumer-year. Policies that expose users to alternative search engines lower Google's market share more than those requiring active choice. We then use Microsoft search logs to assess the impact of additional data on search result relevance. The results suggest that sharing Google's click-and-query data with Microsoft may have a limited effect on market shares.

Hunt Allcott
Stanford University
Stanford, CA 94305
and NBER
allcott@stanford.edu

Juan Camilo Castillo
Department of Economics
University of Pennsylvania
133 South 36th Street
Philadelphia, PA 19104
and NBER
jccast@upenn.edu

Matthew Gentzkow
Department of Economics
Stanford University
579 Jane Stanford Way
Stanford, CA 94305
and NBER
gentzkow@stanford.edu

Leon Musolff
The Wharton School
University of Pennsylvania
3rd Floor, Dinan Hall
3733 Spruce Street
Philadelphia, PA 19104-6372
leon.musolff.1993@googlemail.com

Tobias Salz
Department of Economics, E52-460
MIT
Cambridge, MA 02139
and NBER
tsalz@mit.edu

A randomized controlled trials registry entry is available at
https://www.socialscienceregistry.org/trials/12884

1 Introduction

Search engines are the gateway to the internet, the starting point for 69 percent of all online activities and 44 percent of online purchases.[1] According to a bipartisan report by The US House Judiciary Subcommittee on Antitrust (2020), search engines are part of "the infrastructure of the digital age" and have the potential to "pick winners and losers throughout [the] economy." Due to this key position in the online ecosystem, better search engines can unlock substantial benefits for consumers and firms. Strengthening competition and improving efficiency in the web search market are therefore important policy goals.

Google holds approximately 90 percent of the global web search market (StatCounter, 2024c). Antitrust authorities allege that Google has cemented a dominant market position through anticompetitive practices such as contracts that make Google the default search engine, and that Google's advantages are reinforced by economies of scale in data.[2] Following this line of reasoning, the US Department of Justice sued Google in 2020 (Department of Justice, 2020), and in 2024 the D.C. district court determined in a landmark ruling that Google is a monopolist that has engaged in anticompetitive behavior (D.C. District Court, 2024). Google, however, maintains that its success is driven by its high quality, that competition is "only a click away" given the ease of switching (Page, 2012), and that increasing returns to data are small over the relevant range (Varian, 2015).

We study two questions at the core of this debate. First, why is Google's market share so large? It may be due simply to higher quality. It might stem from users' lack of exposure to alternative search engines, which leads to misperceptions about their quality. It might also result from default effects, which could strengthen Google's position directly through switching costs and inattention, and indirectly by limiting exposure to alternatives. Economies of scale in data could reinforce many of these advantages. Second, how would widely discussed policy interventions such as active choice screens, alternative defaults, and mandatory data sharing impact the market?

To answer these questions, we first develop a model of demand for search engines. Based on our model, we design and implement a field experiment that allows us to identify the sources of Google's market power and estimate our model parameters. Using internal Bing search data, we then estimate economies of scale in data. Finally, we simulate counterfactuals to evaluate the impacts of proposed policy interventions. Throughout our analysis, we focus on the US desktop search market and the competition between Google and Bing, which account for 93 percent of that market (StatCounter, 2024c).

In our demand model, internet users make a binary choice between two search engines (Google and Bing) in each period. In an initial period, their web browser determines the initial default search engine. There are two sources of inertia in switching away from the default: a switching cost that reduces switching by marginal users, and inattention that prevents some users from switching even if they would strongly

[1] See SEO Statistics (2024) and Ecommercedb (2024).

[2] The UK Competition and Markets Authority (2020) has expressed concerns that "we are currently in a catch 22 situation, whereby demand-side remedies would not be sufficiently effective until search engines have access to the level of search data needed to improve their results." The EU's competition authority brought a case against Google in 2003, which resulted in a $4.1 billion fine for Google and the implementation of a choice screen (European Commission, 2018).

prefer to. Users may also begin with incorrect beliefs about the quality of the other search engine. Users switch away from the default if they are paying attention and if the perceived quality improvement of the other search engine outweighs the switching cost.

We designed our experiment both to provide model-free evidence on the key forces and to identify the model's parameters. We recruited a sample of 2,354 desktop internet users from Prolific, a high-quality online survey panel, and asked them to take a baseline survey ("Survey 1") and a follow-up survey ("Survey 2") two weeks later. On Survey 1, participants answered questions about their web search preferences and installed a web browser extension. The extension recorded search engine queries and clicks from two weeks before the first survey until two months after. We also conducted an exit survey at the end of the study.

We randomized participants into a control group and three treatment groups. The Active Choice group was asked what search engine they would like to be their default and then was guided step-by-step through implementing their choice. The Default Change group was offered \$10 to change their default search engine for two days; users who accepted were similarly guided through implementing their choice and then received no further instructions or incentives. The Switch Bonus group was offered a payment to change their default search engine for 14 days, and then asked to make an active choice on Survey 2. The majority of this group was offered a \$10 payment; smaller subsets were offered either \$1 or \$25. Google users in the \$10 Switch Bonus group were further randomized into two interventions implemented by our browser extension: (i) the Ranking Degradation condition, which decreased the relevance of Bing's results by reversing the order of organic search results on the first page, and (ii) the Ad Blocking condition, which removed most ads from Bing's search result pages.

Before the experiment, 96 percent of participants used Google for the majority of their searches. As expected, the Control intervention did not materially affect market shares. The Active Choice intervention also had almost no effect on Google users, of whom only 1.1 percent chose to switch to Bing. The small switching rate suggests that eliminating switching costs and inattention would not meaningfully reduce Google's market share. By contrast, among Bing users, 16 percent chose to switch to Google when required to make an active choice, suggesting that switching costs and inattention have a larger impact on Bing users.

In the \$10 Switch Bonus group, 58 percent of Google users switched to Bing in exchange for our payment. Exposure to Bing increased users' self-reported perceptions of its quality by 0.6 standard deviations. Of those who switched to Bing, 33 percent actively chose to keep using Bing after Survey 2. Our exit survey confirms the importance of learning about Bing, either about its quality or about how to use it: 64 percent of participants who actively decided to keep using Bing reported that it was better than expected, and 59 percent reported that they had gotten accustomed to using Bing. These answers suggest that Google users' lack of experience with Bing is a significant driver of Google's large share at baseline in our sample. For Bing users, the discrepancy between choices before and after exposure to Google is lower, indicating that Bing users are generally well-informed about Google's quality. The survey results again support this interpretation: updating about Google is less pronounced and mostly statistically insignificant.

The Default Change intervention caused an increase in Bing's market share among baseline Google users that persisted until the end of our experiment. Of the users in this group, 81 percent accepted the \$10 to switch to Bing for two days. Among those who complied, we observe a gradual decline in Bing's market share from the initial 100 percent, to 66 percent after one week, 60 percent after two weeks and 46 percent after two months. Through the lens of our model, these patterns suggest that Default Change group participants who keep using Bing do so for two reasons. First, like Switch Bonus group participants, their valuation of Bing increases due to experience. Second, some participants may continue to prefer Google but not switch back due to persistent inattention. Our exit survey is consistent with these hypotheses: 35 percent of Default Change group users who kept using Bing report doing so because they prefer it, while 44 percent report that they forgot to switch back or were too lazy to do so. This result has two implications. First, defaults create a lasting mismatch between preferences and choices. Second, changing defaults can induce learning about unobserved product quality, leading to lasting effects by altering users' perceptions.

Our price treatments uncover substantial heterogeneity in participants' willingness to accept switching search engines. Among Google users, a \$1 payment to switch to Bing for two weeks raised Bing's market share to 32 percent, meaning many users are close to indifferent. With a \$10 payment, Bing's market share increased to 64 percent. At \$25, Bing's market share only increased to 74 percent, meaning many users have strong preferences for Google.

The results from our Ranking Degradation intervention, although somewhat noisy, suggest at most a moderate demand response to search result relevance. While this intervention substantially affected perceptions– it significantly reduced the reported quality of search result relevance and overall search engine quality – we do not detect statistically significant changes in market shares. According to our point estimates, Ranking Degradation reduced Bing's market share by 3.4 percentage points (standard error=2.9). Internal experiments at Google similarly suggest that "a significant quality depreciation by Google would not result in a significant loss of revenues" (D.C. District Court, 2024). However, given our standard errors, our results are consistent with a reduction of up to 9 percentage points.

We then use the data from our experiment to estimate our model by the generalized method of moments. We find that switching costs are negligible, but 34 percent of users are persistently inattentive. The median Chrome user would have to be paid \$3.06 to use Bing instead of Google for two weeks. After two weeks of experience with Bing, the required payment shrinks to \$2.80. Although this dollar difference is small, many users have weak preferences over search engines, so learning shifts market shares significantly. Parameters capturing the responsiveness to search result quality and ad-loads are statistically indistinguishable from zero.

We complement our demand side with estimates of the returns to scale in data using internal search logs from Bing. Specifically, we estimate how click-through rates (a standard measure of result relevance) improve over the life of previously unseen novel queries as Bing serves more results and collects more data. We estimate that returns to data are positive but diminishing, with an approximately logarithmic relationship between cumulative queries and the resulting click-through rate. This relationship predicts that if Bing had access to Google's data, click-through rates would increase from 23.5 percent to 24.8 percent. While there are many potential sources of economies of scale in web search, such as in web indexing and advertising, this analysis of how click-through rates improve with data specifically isolates the benefits relevant for proposed antitrust remedies: those from more click-and-query data. However, this analysis requires strong assumptions, and is therefore more speculative than the rest of our findings.

We consider several counterfactuals. First, we simulate the effect of shutting down all demand-side frictions—switching costs, inattention, and misperceptions about quality. In our model, Bing's market share increases from 11 percent to 26 percent, implying that these frictions are a substantial barrier that prevents Bing from increasing its market share. Consumer surplus increases by $6 per consumer per year.

Next, we simulate the effects of active choice screens that appear when web browsers are first installed, as currently required in the European Union. To simulate the introduction of choice screens on Chrome, we shut down switching costs and inattention, but we continue to allow users to misperceive quality. Driven by the limited effects of our Active Choice intervention, our model predicts that choice screens would increase Bing's market share by only 1.3 percentage points. This small increase underscores that, although choice screens could eliminate certain forms of friction, they would only have a limited impact because they do nothing to avoid the larger barrier to competition that exists because users misperceive Bing's quality.

These results suggest that if regulators want to significantly impact market shares, they should account for the fact that search engines are experience goods (Nelson, 1970) and consider how interventions would impact consumers' exposure to Bing. We use our model to measure the effects of some interventions that could increase exposure. If Google were prevented from bidding to be the default search engine, Bing could become the default on all browsers. This would increase Bing's market share by 40 percentage points. However, it would decrease consumer surplus by $70.92 per consumer per year, because a large number of users will now use Bing even though they strongly prefer Google.

Our results so far highlight a conundrum for competition policy. While choice screens increase consumer surplus by a modest amount, they have almost no effect on market shares. Changing defaults, on the other hand, has a large effect on market shares, but only at the expense of a large decrease in consumer surplus. This raises the question of whether a policy exists that can reduce market concentration without lowering consumer surplus. One possible approach is to mandate that a non-dominant firm—in this case, Bing—be set as the default on all browsers upon installation, followed by a requirement that browsers present a choice screen *after some time*. This would allow users to experience Bing before making an active choice. Such a policy would reduce Google's market share by 16.7 percentage points, while leaving consumer surplus essentially unchanged. Thus, a delayed choice screen could avoid the harm caused by simply setting Bing as the default while reaping the (un-modeled) potential benefits from a less concentrated market, such as increased investment incentives and fewer harms on the advertising side.

Finally, in a more speculative analysis, we account for the feedback effects from endogenous result quality and we simulate the effects of providing Google's search results and click data ("click-and-query" data) to Bing, using the economies of scale estimated from the Bing search logs. Neither feedback effects nor data sharing substantively affect market shares or consumer surplus. This follows from two earlier results: data sharing has only a small effect on Bing's result quality in the Bing search data, and Bing's result quality has a small effect on market shares in the experiment.

What do our results imply for the discussion surrounding Google's dominance? First, we find that most people *do* prefer Google over Bing, and some strongly. At the same time, Google also significantly benefits from frictions that raise its market share beyond the efficient level. The debate has largely focused on the $26 billion that Google spends annually to secure its default position on browsers and Android devices (D.C. District Court, 2024), and our findings confirm that defaults play an important role. However, our results suggest that the power of Google's default position on Chrome does not stem from preventing users from choosing Bing, since most consumers at least think they prefer Google. Instead, Google's default position ensures that users are never exposed to Bing, and hence never learn about it. Such learning would permanently lower Google's market share in our model. Our results suggest that regulators and antitrust authorities can increase market efficiency by considering search engines as experience goods and designing remedies that induce learning. This conclusion may be of broader relevance. Prior theoretical literature has shown that incumbents might benefit from favorable user expectations (Schmalensee, 1982). Our work confirms this fact and its importance empirically.

Our results have several important limitations. First, we focus only on desktop search in browsers because parts of our experiment cannot be implemented on mobile. Desktop is important *per se*, representing 55 percent of the total (search + non-search) web traffic in 2023 (StatCounter 2024b), but switching costs could be higher on mobile and on non-browser search bars integrated into the Windows, Android, and iOS operating systems. Second, our experiment sample is more educated, has higher income, and is more white than the population of US adults, and it may also not be representative on unobserved factors such as price sensitivity or computer literacy. Third, our economy of scale analysis requires strong identifying assumptions.

Our work contributes to several related literatures. First, we contribute to the literature on competition and antitrust concerns in the web search market (Ostrovsky, 2021; Vásquez Duque, 2022; Decarolis, Li and Paternollo, 2023; Hovenkamp, 2024) and the surrounding policy discussion (Patterson, 2012; Stigler Committee on Digital Platforms, 2019; UK Competition and Markets Authority, 2020; Dinielli et al., 2023; Heidhues et al., 2023).[3] Two papers are particularly related to our work. Decarolis, Li and Paternollo (2023) uses observational data to investigate the effect of antitrust remedies imposed by European and Russian regulators. They find small effects from introducing a choice screen in the EU, consistent with the results from our Active Choice treatment. Our work goes beyond theirs by offering an explanation for the small effects of choice screens, more broadly investigating the sources of Google's large market share, and quantifying the equilibrium effects of different remedies. Vásquez Duque (2022) conducts a survey experiment on Amazon Mechanical Turk in which participants choose search engines under two conditions: an active choice treatment and a default treatment. His findings suggest that active choice has a

[3] Most existing work on the search engine market has focused on the advertising side (Varian, 2007; Edelman et al., 2007; Athey and Ellison, 2011; Blake et al., 2015). In addition, there are studies assessing the value of digital services that don't charge prices to consumers (Brynjolfsson et al., 2019), of which search engines are an important example.

small effect on market shares, and that misperceptions may significantly contribute to Google's high market share. Our work goes beyond Vásquez Duque in three main ways: we conduct a field experiment based on incentivized real-world choices, our browser extension allows us to implement additional treatments essential for disentangling the sources of Google's market power, and we model counterfactuals that speak directly to policy.

Second, we extend previous work on the competitive effect of choice frictions, including the effect of switching costs (Klemperer, 1987; Farrell and Klemperer, 2007) and the importance of defaults in the presence of inattention (DellaVigna and Malmendier, 2006; Carroll et al., 2009; Handel, 2013; Ericson, 2014; Ho et al., 2017; Andersen et al., 2020; Fowlie et al., 2021; Einav et al., 2023; Miller et al., 2023; Brot-Goldberg et al., 2023; Lee and Musolff, 2023). Our results highlight an important new role of defaults in a market setting. Even when switching costs and inattention are relatively small, defaults can matter by preventing consumers from gaining experience with an alternative whose quality they initially underestimate. Like Agte et al. (2024), we thus find that switching costs and misspecified beliefs interact with each other.

Third, we extend the empirical literature that studies experience goods (Erdem and Keane, 1996; Ackerberg, 2003; Israel, 2005; Crawford and Shum, 2005; Dickstein, 2021) by showing that overly pessimistic consumer beliefs about the quality of rivals help entrench dominant firms.

Fourth, we extend previous empirical work on economies of scale in search (Chiou and Tucker 2017; He et al. 2017; Azevedo et al. 2020; Schaefer and Sapi 2023; Klein et al. 2023), and in data more broadly (Bajari et al. 2019; Tucker 2019), by combining our estimates of the returns to scale with experimental estimates to quantify the equilibrium implications of antitrust remedies and the resulting welfare effects for consumers.

Fifth, and more broadly, we contribute to a literature that experimentally studies digital markets. Prior studies have focused on consumer surplus from social media (Allcott et al., 2020), price salience (Blake et al., 2021), addiction to digital services (Allcott et al., 2022), substitution pattern across online services (Aridor, 2022), and the welfare consequences of platforms when people experience fear of missing out (Bursztyn et al., 2023). To facilitate such studies, Farronato, Fradkin and Karr (2024) introduced an open source browser extension, whose code was helpful in developing the extension for this project.

Sections 2–8, respectively, present the model, experimental design, data, model-free experimental results, structural estimation, economy of scale analyses, counterfactuals, and conclusion.

2 Model

We now present our model of demand for search engines, which guides our experimental design. The experiment's results will be used to estimate the model's parameters and explore the effect of different antitrust policies.

2.1 Search engine choices

Consumers indexed by i make a binary choice between two search engines $j \in \{B, G\}$ in two-week periods indexed by t. We think of this choice as determining the search engine for both direct navigation and address bar searches.[4] Consumer i's search engine choice in period t is denoted y_{it}.

Each consumer has an exogenously set web browser (Chrome or Microsoft Edge), which determines her default search engine, i.e., the search engine used for address bar searches when a browser is first installed. We denote the default search engine by d, and the alternative search engine by $-d$. For Chrome users, the default is Google ($d = G$), and the alternative search engine is Bing ($-d = B$). For Microsoft Edge users, the default is Bing and the alternative search engine is Google.

The price agent i receives for using search engine j in period t is p_{ijt}. Normally $p_{ijt} = 0$, but our experimental interventions will vary prices. Variables a_j^*, r_j^*, and ξ_j^* refer to j's ad load, search result relevance, and other unobserved characteristics respectively. We define $\zeta_j^* := \alpha a_j^* + \rho r_j^* + \xi_j^*$ as j's "quality." The term χ_{ij} is an idiosyncratic preference shifter that does not vary across time periods.

Each period's flow utility from j is

$$u_{ijt}^* = \eta p_{ijt} + \zeta_j^* + \chi_{ij}, \tag{1}$$

with $\eta > 0$.

Users may be imperfectly informed about the quality of search engines. We use $\mathbb{E}_{it}$ to denote agent i's expectation of different quantities at time t. Thus, the quality agent i perceives at time t is given by $\mathbb{E}_{it}[\zeta_j] := \alpha \mathbb{E}_{it}[a_j] + \rho \mathbb{E}_{it}[r_j] + \mathbb{E}_{it}[\xi_j]$. Although the true quality ζ_j^* is constant, the perceived quality $\mathbb{E}_{it}[\zeta_j]$ depends on time because users' perceptions may change over time, as we explain below. The perceived flow utility is given by

$$u_{ijt} = \eta p_{ijt} + \mathbb{E}_{it}[\zeta_j] + \chi_{ij}. \tag{2}$$

We assume that users have correct beliefs about their default search engine d, so $\mathbb{E}_{it}[\zeta_d] = \zeta_d^*$ and $u_{idt} = u_{idt}^*$. Users that have never chosen the alternative search engine $-d$, on the other hand, may be imperfectly informed about its quality. In that case, perceived quality $\mathbb{E}_{it}[\zeta_{-d}]$ takes a different value $\tilde{\zeta}_{-d} := \alpha \tilde{a}_{-d} + \rho \tilde{r}_{-d} + \tilde{\xi}_{-d}$. After one period of experience with $-d$, consumers become fully informed, so their perceived quality becomes $\mathbb{E}_{it}[\zeta_{-d}] = \zeta_{-d}^*$.

There are two sources of inertia. First, there is a switching cost σ of getting to the choice screen and changing the default. Second, consumers may be inattentive. A fraction ϕ of users never pay attention (we say they are "permanently inattentive"), so they always stay with their default search engine. The remaining fraction $1 - \phi$ of users are probabilistically inattentive. In each period, with exogenous probability π they are attentive and thus consider the choice between search engines, and with probability $1 - \pi$ they are inattentive.

[4] In our experimental sample, only 6.6 percent of users search more than 10 percent of the time on a search engine that is not their current browser default.

In period t, if inattentive, consumers stick with the search engine they used in the previous period: $y_{it} = y_{i,t-1}$. If attentive, the consumer chooses the search engine that maximizes utility over an infinite horizon with per-period discount factor δ:

$$y_{it} = \arg\max_{j\in\{G,B\}} \left\{ u_{ijt} - \sigma \mathbf{1}_{j\neq y_{i,t-1}} + \delta V_{i,t+1}(j) \right\}. \tag{3}$$

where $V_{i,t+1}(j)$ is the perceived continuation value after having chosen search engine j.

We make three assumptions that simplify this dynamic switching problem into an effectively static decision. First, consumers do not perceive any uncertainty about quality ζ_{-d}, so there is no option value to exploration. Second, consumers weakly underestimate the quality of the alternative search engine ($\tilde{\zeta}_d \leq \zeta^*_{-d}$), so experience with $-d$ weakly increases its market share. Finally, we assume that by the start of the experiment all participants have made an attentive choice, so that market shares are in steady state at $t = 0$. This approximates a world in which the time between browser installation and the beginning of the experiment is long.

We believe that the first assumption is psychologically realistic, and also consistent with what we observe in the Active Choice treatment. The second assumption is also consistent with our experimental estimates of experience effects. The final assumption is consistent with the lack of market share trends before our experiment and in our Control group.

Given these assumptions, and since idiosyncratic preferences remain constant over time, the consumer's decision is effectively static: if it is optimal to switch in the future, it is optimal to switch immediately. Thus, attentive consumers permanently choose either B or G, where they account for the perceived discounted utility $u_{ijt}/(1-\delta)$ of each search engine. Equation (3) thus simplifies to

$$y_{it} = \arg\max_{j\in\{G,B\}} \left\{ \frac{u_{ijt}}{1-\delta} - \sigma \mathbf{1}_{j\neq y_{i,t-1}} \right\}. \tag{4}$$

At baseline—before the experiment starts—prices are equal to zero.[5] Hence, the perceived discounted utility from permanently continuing with the default search engine is

$$\frac{\zeta^*_d + \chi_{id}}{1-\delta}. \tag{5}$$

The utility from permanently switching to the alternative search engine is

$$\frac{\tilde{\zeta}_{-d} + \chi_{i,-d}}{1-\delta} - \sigma. \tag{6}$$

We now define variables for differences between the alternative search engine and the default search engine: $\Delta v_t := v_{-d} - v_d$ for any variable v. Concretely, $\Delta p_{it} := p_{i,-d,t} - p_{idt}$, $\Delta\zeta^* := \zeta^*_{-d} - \zeta^*_d$, and $\Delta\chi_i := \chi_{i,-d} - \chi_{i,d}$.

[5] In practice, users can earn rewards for Bing searches and redeem them for Microsoft products. However, those rewards are modest so we do not model them directly. This modeling decision effectively means that rewards are captured by the quality term ζ.

We also define $\Delta\tilde{\zeta} := \tilde{\zeta}_{-d} - \zeta^*_d$, since users always perceive the utility of the default search engine correctly. With this notation and after differencing 5 and 6, an attentive consumer switches to $-d$ if

$$\Delta\tilde{\zeta} + \Delta\chi_i - \sigma(1-\delta) > 0, \tag{7}$$

Therefore, the probability that a consumer that pays attention chooses $-d$ is $\mathbb{P}(\Delta\tilde{\zeta} - \sigma(1-\delta) > -\Delta\chi_i)$. Defining $S(\cdot)$ as the cumulative density function of $-\Delta\chi_i$, this becomes $S\left(\Delta\tilde{\zeta} - \sigma(1-\delta)\right)$.

A fraction ϕ of users is permanently inattentive, and thus will keep d forever. The remaining fraction $(1-\phi)$ will eventually pay attention at some point and choose the search engine that maximizes perceived utility. Thus, at time $t = 0$, $-d$'s market share is

$$s_{-d,0} = (1-\phi)S\left(\Delta\tilde{\zeta} - \sigma(1-\delta)\right). \tag{8}$$

Equation (8) illustrates four reasons why Google might have high steady-state market share: (i) Google has a higher true quality than Bing ($\Delta\zeta^* < 0$), (ii) consumers perceive Bing to be worse than it actually is ($\tilde{\zeta}_B < \zeta^*_B$), (iii) Google is many users' initial default and the switching cost σ is large, or (iv) Google is many users' initial default and the fraction ϕ of permanently inattentive users is large. Optimal choices are attained in a counterfactual where consumers learn the true ζ and make active choices with zero switching cost. In that scenario, the market share of $-d$ is $s_{-d,0} = S(\Delta\zeta^*)$.

In Section 5 we show formally how our experimental treatments identify the parameters (η, σ, π, ϕ,$\Delta\zeta^*$, $\tilde{\zeta}_{-d} - \zeta^*_{-d}$, α, and ρ) of this model.

3 Experimental Design

3.1 Overview

Figure 1 illustrates the experimental design and timeline. There are two surveys. Survey 1 took place immediately after recruiting. The invitation to Survey 2 was sent the morning of the 15th day after participants complete Survey 1. The experiment ended two months after Survey 1. We sent an exit survey to some users at the end of the experiment.

Recruitment, screening, and demographics. From March 19th to April 2nd, 2024, we recruited participants from the Prolific online platform, enforcing balance by gender. To qualify for the study, participants had to be US residents at least 18 years old.

Survey 1 began with screening questions concerning the participant's device, web browser, and search engine use. To obtain a survey-based measure of their current address bar search engine, we asked participants to search for the term "potato" through their browser's address bar and report the search engine that they were directed to. Unless users intentionally changed the address bar search engine, it is the browser's

Figure 1: **Experimental Design**

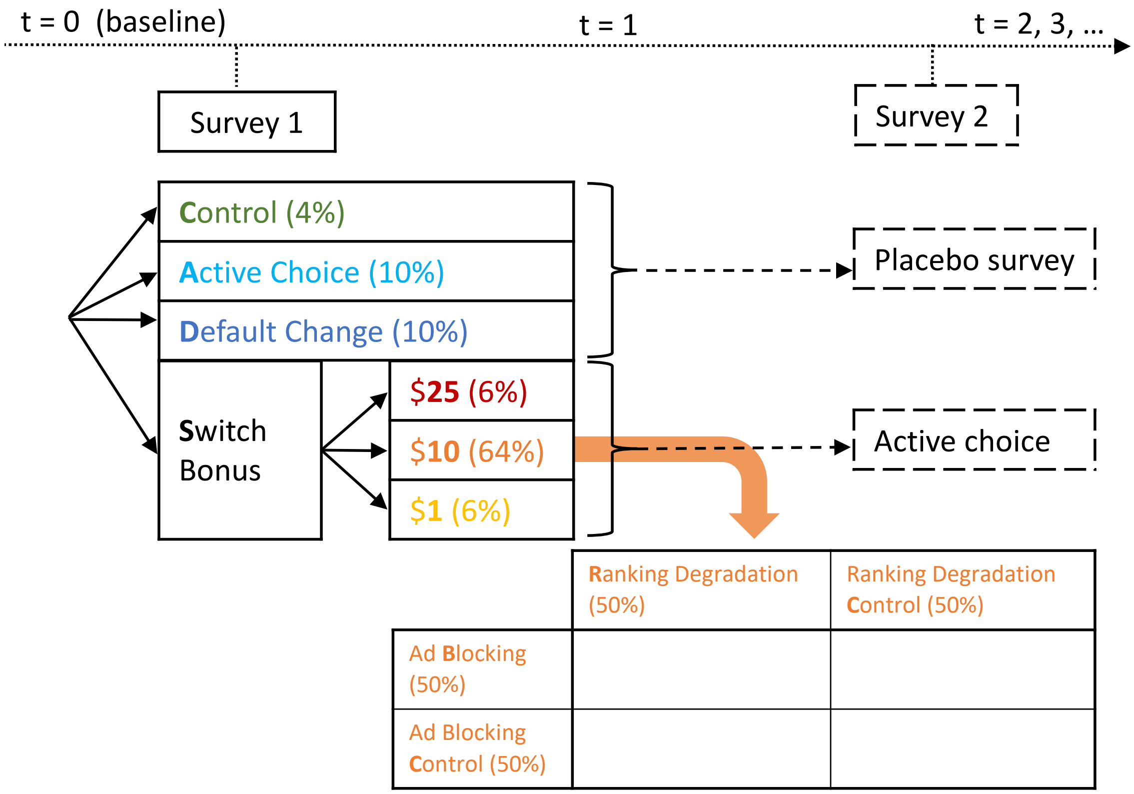

Notes: This figure illustrates our experimental design. The Control group was guided through how to change their bookmarks. The Default Change group was paid $10 to switch their default from Google to Bing for two days. The Active Choice group was asked about their preferred default and then guided to implement this choice. The Switch Bonus treatment group was asked to switch search engines in return for a bonus payment of either $1, $10, or $25. The $10-group was further randomized in a two-by-two factorial design, which varies whether ads were blocked and whether search results are degraded. The Switch Bonus group made an active choice after the 14-day incentive period.

default search engine. The survey then separately asked, "What search engine do you usually use on this web browser?" in case users do not usually search via the address bar.

Participants could continue with Survey 1 only if they accessed the survey through a desktop or laptop computer using either Edge or Chrome and they reported that (i) on their current device, they exclusively use the current web browser, (ii) they do not frequently share that computer with other people, (iii) their address bar search engine is either Google or Bing, and (iv) the search engine they usually use is either Google or Bing We restrict the analysis to participants with a "consistent baseline search engine": the search engine (Bing or Google) they report they usually use is also (i) the search engine they reported as the address bar search engine (ii) the search engine for more than half of searches recorded by our browser extension before installation, and (iii) the address bar search engine recorded by our browser extension. We also drop users with fewer than 10 recorded searches in the 20 days before installing our browser extension during Survey

1.

Participants who passed the screening questions and consented to participate were then asked demographic questions. This was followed by a series of questions eliciting opinions about Google and Bing, including why they use the search engine they usually use.

Search engine rating questions. For all participants on Survey 1 and some participants on Survey 2, we asked four *search engine rating questions*. We first asked people to rate Google versus Bing in terms of overall quality and then on specific dimensions. Possible answers were *Google is a lot better*, *Google is a little better*, *they are about the same*, *Bing is a little better*, and *Bing is a lot better*. The answer order was randomly flipped, so that half of participants saw "Google is a lot better" on the left and half on the right.

We asked about the following quality dimensions: (i) relevance and ordering of search result links, (ii) features on search result pages (e.g., weather info), (iii) relevance of ads, (iv) AI chat, (v) privacy, and (vi) rewards or loyalty points. The possible answers were the same as before. We also included an attention check, to "please choose 'Bing is a lot better' if you are still paying attention." The rows were presented in random order.

Search Extension. Participants were then asked to install Search Extension, a browser extension developed for this study. Figure A1 shows what users see when they first install the extension. As a standard Chrome/Edge browser extension, it is unobtrusive and not visible on the browser interface after the installation. Search Extension records the dates, times, and information identifying the source (the address bar or the search engine website) of all searches on all general web search engines (google.com, bing.com, etc.) that take place after installation. Using the browser's recorded search history, it also collects the same information for all searches made in the 20 days before installation.[6] Additionally, for searches made after installation, the extension records whether the user clicked on an ad or an organic search result, and if so, the rank of the result.

Search Extension includes two intervention functionalities that we turned on or off in treatment conditions described below. First, the Ranking Degradation functionality reverses the order of organic results on search result pages. Thus, the bottom results are moved to the top, and the top results are moved to the bottom. Second, the Ad Blocking functionality removes all ads that it detects on search result pages. Search Extension does not make users aware of these functionalities or whether they are turned on. These interventions occur at a split-second when the page loads and are imperceptible to the user.

Compensation. In addition to the incentive payments associated with each treatment, participants were paid a base payment of $25: $5 each for completing Survey 1 and Survey 2, $5 for installing Search Extension, and $10 for keeping Search Extension installed for two months after completing Survey 1.

[6] Throughout our analysis, we only use data for the two weeks prior to installation to harmonize the data with our experiment.

3.2 Treatment Groups

Users were randomized into four treatment groups, one of which has further sub-treatments. Participants whose baseline default search engine was Google were randomized into all groups, with the proportions in Figure 1 and below. Participants whose baseline default search engine was Bing were randomized into only two groups (A and S10CC, see below), with 50-50 probability. We randomized Bing users only into these two groups to increase power, because there are relatively few Bing users. We did not include a Control group for Bing users because market shares are stable over time without intervention.

We now describe each group's experience after installing Search Extension on Survey 1.

Control (group "C," 4 percent of baseline Google users). The Control group was shown information about how to change the bookmarks on their web browser, in a similar format to the treatment information the other groups receive. We correctly anticipated that this placebo intervention would not change search engine market shares.

Active Choice (group "A," 10 percent of baseline Google users). The Active Choice group was told that we would now show them how to change the default search engine. To avoid experimenter demand effects, the survey clearly stated that "whether you change it or not is up to you." The survey then asked, "when we get to the screen where you can set your default search engine, what would you like your default to be?" We call this the person's desired default. The survey then showed people how to change the default search engine, asked people to copy and paste the correct settings page URL (to confirm that they were on the correct page), asked people to set the address bar default search engine to their desired default, and asked to confirm that they had done so or explain why not.

Default Change (group "D," 10 percent of baseline Google users). The Default Change group was told that we would pay them $10 if they switch their default to Bing and make at least 90 percent of their searches (and at least 4 searches in total) on Bing over the next 2 days. The survey told participants that we would show them how to change the default search engine, and then asked, "Would you like to accept the additional $10 to make Bing your primary search engine for the next 2 days?" As in the Active Choice condition, the survey then showed people how to change the default search engine and asked people to copy and paste the correct settings page URL. For participants that said they would like to accept the offer, the survey asked people to change the address bar default search engine and to either confirm that they have done so or explain why not.

Our ideal version of this intervention would be to just change the default search engine automatically, without telling participants how to change it. Browser extensions are not able to change default settings on Chrome or Edge, so we were unable to do this. The experience described above approximates such an intervention.

For these first three groups (C, A, and D), we wanted a second survey only to make sure that all groups have the same number of surveys. Thus, Survey 2 for those three groups is a "placebo survey."

Switch Bonus (group "S," 72 percent of baseline Google users). Following the same steps as for the Default Change groupl, the Switch Bonus group was offered payment in return for switching search engines for fourteen days and for making at least 90 percent of their searches (and at least 20 searches in total) on the alternative search engine during this time. Payments p were experimentally varied between \$1, \$10, or \$25, with 6, 64, and 6 percent probability, respectively. To avoid considerations of future inertia, participants were told that "on the second survey in 14 days, we will remind you how to switch your default search engine."

On Survey 2, the 14-Day Switch group was asked the search engine rating questions and then received the active choice intervention described above.

For baseline Google users, the \$10 Switch Bonus group (with 64 percent of the sample) was further factorialized into a two-by-two matrix of two search results interventions implemented by Search Extension:

- **Ranking Degradation** (group "R," 50 percent of S10 group). Upon installation, Search Extension turned on the ranking degradation functionality on Bing.

- **Ad Blocking** (group "A," 50 percent of S10 group). Upon installation, Search Extension turned on the ad blocking functionality on Bing.

We refer to the 25 percent subset of the S group assigned to the Ranking Degradation Control and Ad Blocking Control groups as group S10CC. We made the \$10 Switch Bonus group relatively large and the Ad Blocking and Ranking Degradation conditions relatively forceful because we expected limited power to detect the effects of these two interventions on market shares.

3.3 Exit Survey

At the end of the experiment (two months after Survey 1), we sent an exit survey to a random subset of participants in the Switch Bonus and Default Change groups whose original search engine was Google and who kept Bing after the incentive period. Eligible participants were offered \$5 for completing it. The survey experience was similar across the Switch Bonus and Default Change groups, each of which was first asked "Over the past 6 weeks, our records show that you have continued to use Bing on this browser. Why?" with a free-form text response field. Subsequently, they were asked "Why did you decide to keep using Bing? Please choose all that apply." The available options were different for both groups.

For the Switch Bonus group, the possible answers were (i) *Before the experiment, I had always wanted to use Bing but hadn't gotten around to it,* (ii) *Bing was better than I thought it would be,* (iii) *I got accustomed to Bing,* and (iv) *Other*, with forced free form response.

For the Default Change group, the possible answers were (i) *I wanted to keep using Bing,* (ii) *I forgot to change back to Google, (iii) Changing back to Google was too much effort,* and (iv) *Other*, with forced free form response.

3.4 Pre-Analysis Plan

We submitted our final pre-analysis plan (PAP) to the AEA RCT registry in February 2024, the month before data collection began.[7] The PAP specified how we would construct the basic experimental results, by presenting mockups of Tables 1, 2, 5, A1, and A2, and Figures 1, A3, 2, 3, and 5. Our tables and figures follow the PAP's mockups except for slight reformatting and the addition of one row in Table 1.

4 Experimental Results

4.1 Data

Table 1 shows the sample sizes at each point of the recruitment and experiment. The sample reported in each row is a strict subset of the row above. There were 45,219 Prolific users in the US eligible to participate in the main experiment. Of those, 5,280 people saw the study advertisement on Prolific and 4,217 started Survey 1, of whom 2,736 passed our screening questions and consented, 2,518 completed Survey 1. Of those 2,518, 128 people were disqualified because they tried to retake the survey after being disqualified in their first attempt, 36 failed to correctly install Search Extension, 440 had fewer than 10 recorded baseline searches, and 188 did not have a consistent baseline search engine, leaving an initial sample of 1,726 participants. Of those, 1,660 finished Survey 2, and 1,461 kept Search Extension installed until the end of the experiment (two months after Survey 1). Those 1,461 people form our final analysis sample for all results tables and figures.

Table 2 shows demographic covariates for the sample and for US adults. Our sample is nearly balanced on gender but younger, more educated, more white, and higher-earning than the US adult population. It also includes a higher share of desktop Google users.

In the Appendix, we present tests of balance and differential attrition. Appendix Table A1 shows that, for both baseline Google and baseline Bing users, treatment assignment is statistically balanced on the covariates reported in Table 2. Appendix Table A2 presents our tests for differential attrition across treatment groups at various stages of the experiment. Among Google users, we reject the null hypothesis of balanced attrition after Survey 2. Surprisingly, this is mostly driven by the S25 group, who can earn the largest incentive payment. Among users assigned to the S25 group, 92 percent finish Survey 2, whereas completion rates in the other treatments range from 92 percent to 99 percent. However, importantly, we cannot reject the null hypothesis of balanced deinstallation of the extension at the end of the experiment (even between S25 and other groups). Depending on the treatment, we collect extension data from 93 percent to 97 percent of

[7] See https://www.socialscienceregistry.org/trials/12884.

Table 1: **Sample Sizes**

	Sample size
US Prolific users not in pilots	45,219
Saw study advertisement	5,280
Started Survey 1	4,217
Passed screening questions	2,737
Consented	2,736
Finished Survey 1	2,518
Not rejected for multiple responses	2,390
Installed Search Extension	2,354
At least 10 baseline searches	1,914
Consistent baseline search engine	1,726
Finished Survey 2	1,660
Kept Search Extension 2 weeks after Survey 2	1,577
Kept Search Extension 2 months after Survey 1	1,461

Notes: This table presents sample sizes at each stage of the experiment. "Not rejected for multiple responses" drops users who tried to retake the survey after being disqualified in their first attempt. "Consistent baseline search engine" requires that the same search engine (Bing or Google) is (i) reported on Survey 1 as the search engine used for address bar search, (ii) reported as the search engine they usually use, (iii) the search engine for more than half of searches recorded by Search Extension before installation, and (iv) the address bar default recorded by Search Extension before installation. "Kept Search Extension 2 weeks after Survey 2" refers to the users whose Search Extension was active 2 weeks after their Survey 2 was completed. "Kept Search Extension 2 months after Survey 1" refers to the users whose Search Extension was active 8 weeks after the beginning of the experiment. The sample in each row is a strict subset of the row above.

Table 2: **Sample Demographics**

	(1) Analysis sample	(2) U.S. adults
Income ($000s)	56.45	40.86
College	0.58	0.33
Male	0.45	0.49
Age	36.39	48.16
White	0.60	0.32
Baseline Google user	0.96	0.82

Notes: Column 1 presents average demographics for our analysis sample, and column 2 presents average demographics of American adults. All but the bottom row in column 2 use data from the 2022 American Community Survey 1-year estimates. There is no nationally representative estimate of the share of US adults that use Google vs. Bing, but StatCounter (2024a) reports that Google and Bing had 76 and 17 percent desktop search market shares in the US in December 2023, meaning that Google had 82 percent of the Google plus Bing market.

Table 3: **Summary Statistics**

Panel A: Search Activity		
	Before Experiment Mean (SE)	*During Experiment* Mean (SE)
Fraction (%) of days with positive search	62.0 (3.5)	65.2 (3.3)
Daily searches	11.105 (1.779)	10.911 (1.159)

Panel B: Pre-Experiment Search Engine Share	
	Share (%) of Google Searches
Google users	97.8 (0.1)
Bing users	3.5 (0.8)

Panel C: Users by Browser and by Pre-Experiment Search Engine (percentages)		
	Google	*Bing*
Chrome	94.3	1.2
Edge	1.2	3.3
Total	95.6	4.4

Notes: This table provides summary statistics of participants' search behavior both before (pre Survey 1) and during the experiment (between Survey 1 and eight weeks after Survey 1). The statistics show the mean and standard error of individual-level averages.

participants two weeks after Survey 2. Starting two weeks after Survey 2 more users uninstall the extension. The attention check embedded in the search engine rating questions is passed by 96 percent of Google users and 100 percent of Bing users.

Table 3 presents summary statistics. Panel A shows that participants conduct at least one search on over 60 percent of days, averaging about 11 searches daily. Using pre-experiment browser history data collected by Search Extension, we compare search behavior before and during the experiment (i.e., the eight weeks after Survey 1). We find no statistically significant difference in either the frequency of search days or the average number of daily searches between these periods, suggesting that any substitution to other browsers or devices is limited.[8] Panel B shows that after limiting to users with a consistent baseline search engine, there is only a minimal amount of multi-homing in our sample for both Google and Bing users.

4.2 Initial Survey Ratings of Google and Bing

Figure 2 shows how participants rate search engine quality on Survey 1. Figure 2 presents the average quality ratings on a five-point scale, coded from -2 (*"Google is a lot better"*) to +2 (*"Bing is a lot better"*).

Google users strongly prefer their search engine on almost all dimensions, including overall quality, the relevance of search results, result page features, relevance of ads, AI chat integration, and privacy. The only

[8] We further investigate the number of daily searches by treatment group in Appendix Table A3 and Appendix Figure A5. The only statistically significant result we find is that those baseline Google users that we are paying $10 to use Bing search more during the incentive period, perhaps driven by the need to conduct more searches to find the right result. If users had substituted away to other devices, we would have found the opposite.

Figure 2: **Initial Ratings of Google and Bing**

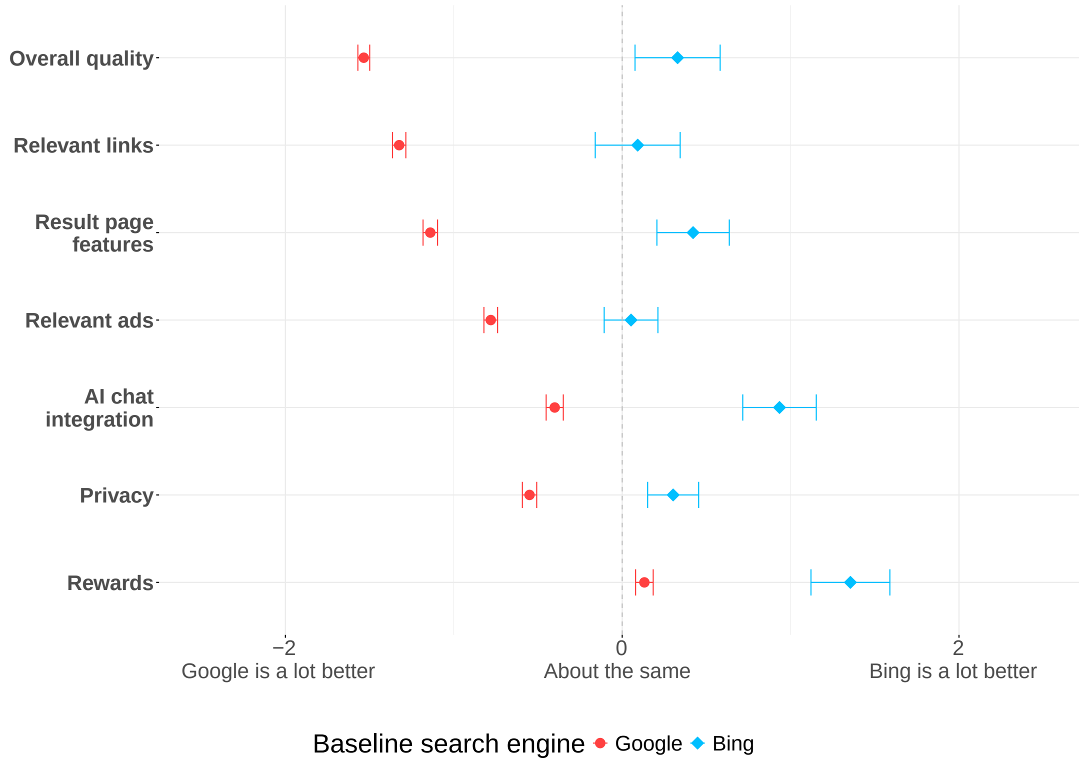

Notes: This figure presents average responses to the search engine rating questions for baseline Google and Bing users on each reported dimension, in response to the following questions: "Overall, how would you rate the quality of Google relative to Bing?" and "How would you rate the quality of Google relative to Bing on the following dimensions?" Response options were "Bing is a lot better," "Bing is a little better," "They are about the same," "Google is a little better," and "Google is a lot better," coded as 2, 1, 0, -1, and -2, respectively. Whiskers indicate 95 percent confidence intervals.

exception is rewards, for which they express a slight preference for Bing. This preference is consistent with the fact that using Bing can earn users Microsoft Rewards worth up to $10 per month, whereas Google offers no such rewards program. Meanwhile, Bing users only slightly prefer Bing on overall quality, result page features, AI chat integration, and privacy, but strongly prefer Bing in terms of rewards. Bing users rate the search engines similarly on relevance of links and ads.

We also asked participants why they chose Google or Bing. A significant share of participants attribute their usage to the browser's default setting. These shares are comparable across Bing (53 percent) and Google users (56 percent). Note that many of these users actually prefer their default search engine, and might have switched to it if they had been given a different default, so their reports need not mean the default is decisive in a causal sense. Bing users are less likely to report results page features (7 and 17 percent respectively) and relevance of links (26 and 69 percent respectively) as reasons for choosing their search engine. By contrast, Bing users are more likely than Google users to report AI integration (29 and 9 percent respectively) and rewards (66 and 2 percent respectively). Appendix Figure A3 presents full results.

4.3 Effects of Main Treatments

Figures 3 and 4 present market shares by treatment group over time. To compute the market share at time t, we first compute the market share for each participant separately and then average shares across participants.[9]

Figure 3 presents results for baseline Google users. The $10 Switch Bonus group presented here is limited to S10CC participants, who did not experience the Ad Blocking or Ranking Degradation interventions. Figure 4 limits to baseline Bing users. It only includes the two groups to which they were assigned, Active Choice and Switch Bonus with Search Extension Intervention Control (S10CC).

Baseline and Control. Figure 3 shows that the Bing market share among baseline Google users was virtually zero before the experiment, reflecting the fact that most of these users exclusively used Google. We also see in panel (a) that control group users did not change their behavior during the experiment: 0.5 percent of overall searches in the fourteen days after Survey 1 were made on Bing. Figure 4 shows that baseline Bing users did use Google occasionally before the experiment, but only for about 4 percent of searches.

Active Choice. Figure 3 shows that the Active Choice intervention had almost no effect on the search engine usage of baseline Google users, increasing Bing's market share from 0.7 percent to just 1.9 percent. This result is consistent with the relatively small effect of the choice screen that Google implemented on Android devices in response to the European Commission's antitrust ruling in 2018 (Decarolis et al., 2023; European Commission, 2018). This means that most baseline Google users choose Google even when they

[9] We can also define market shares by treating search engine choice as binary (defined as the search engine that the person used the most in period t) and then taking the average of those binary choices. This gives very similar results.

Figure 3: **Search Market Shares for Baseline Google Users**

(a) **Active Choice, Control and Default Change**

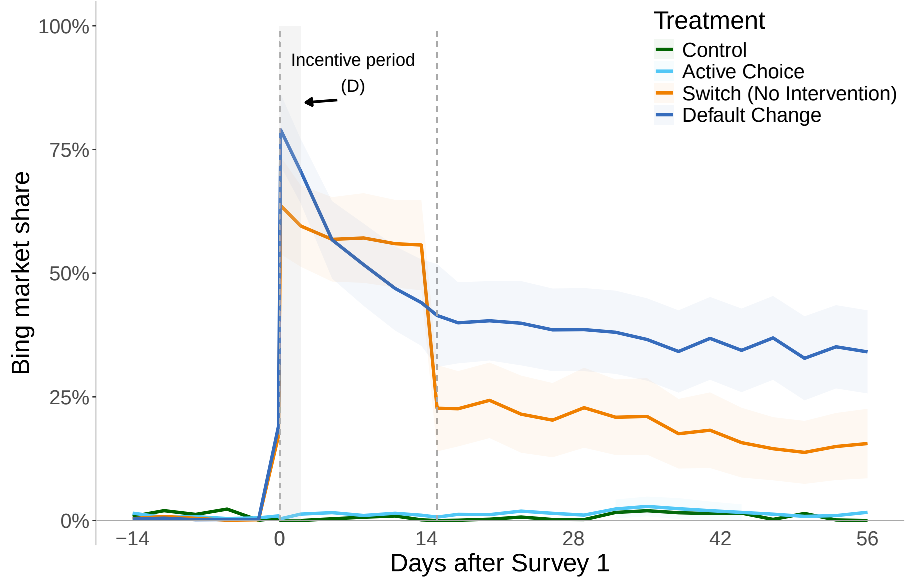

(b) **Switch Bonus**

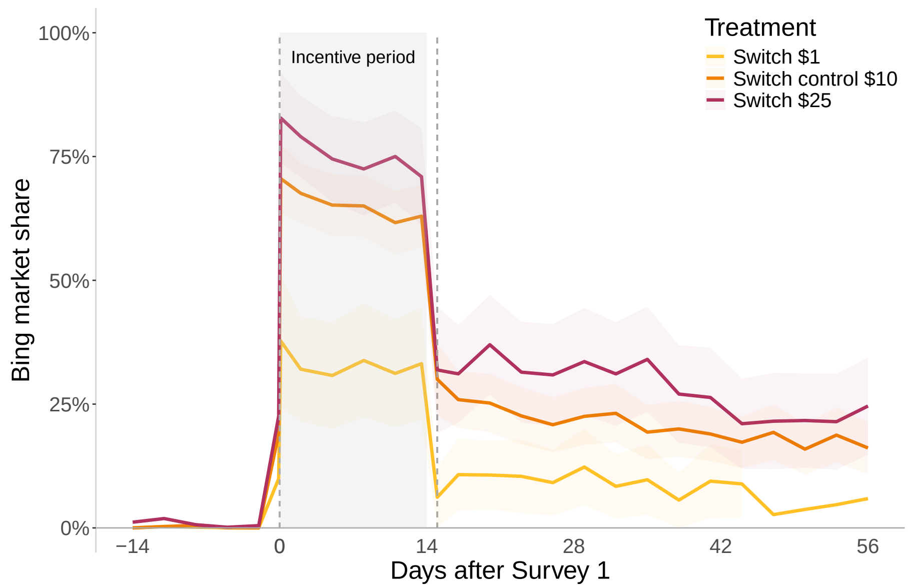

Notes: This figure presents Bing market shares by treatment group for each day of the experiment, for baseline Google users. The dashed vertical lines mark the dates of the two surveys, and the shaded area indicates the incentive period. To compute daily market shares, we first compute the daily market share for each participant and then average shares across participants. To smooth and clarify the figure, we plot averages for groups of two or three days. We define day 15 as the day of Survey 2, so, for any participants who did not take Survey 2, we drop any data after day 14. For any Survey 2-takers who did not take Survey 2 on day 15, we drop any post-day 14 data recorded prior to Survey 2 completion. As we define day 15 as the day of Survey 2, the post-Survey 2 market shares are displayed from day 15 for all Survey 2-takers. In panel 4a, "Switch (No Intervention)" refers to all users in the S1, S10CC, and S25 groups.

Figure 4: **Search Market Shares for Baseline Bing Users**

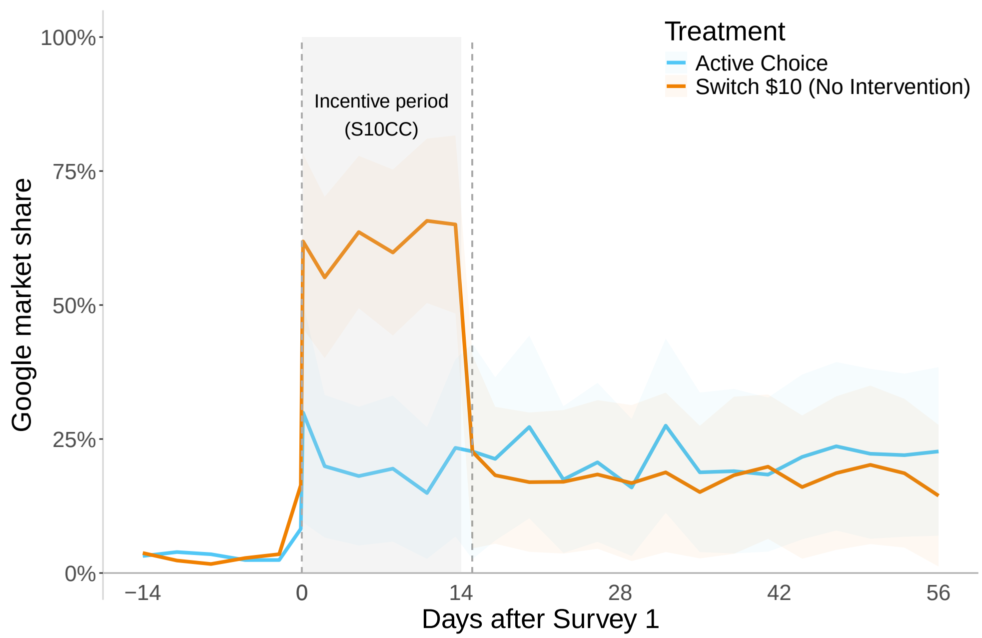

Notes: This figure presents Google market shares by treatment group for each day of the experiment, for baseline Bing users. The dashed vertical lines mark the dates of the two surveys, and the shaded area indicates the incentive period. To compute daily market shares, we first compute the daily market share for each participant and then average shares across participants. To smooth and clarify the figure, we plot averages for groups of two or three days. We define day 15 as the day of Survey 2, so, for any participants who did not take Survey 2, we drop any data after day 14. For any Survey 2-takers who did not take Survey 2 on day 15, we drop any post-day 14 data recorded prior to Survey 2 completion. As we define day 15 as the day of Survey 2, the post-Survey 2 market shares are displayed from day 15 for all Survey 2-takers.

are attentive and there are no switching costs, indicating that removing these frictions alone would not change Google's large market share.

Among baseline Bing users, on the other hand, we find that the Google market share increased from 3.8 percent to 18 percent in the Active Choice group. The larger effect for Bing users is consistent with many permanently inattentive Edge users preferring Google—and thus switching when given an active choice—but relatively few permanently inattentive Chrome users preferring Bing. It could also reflect switching costs being larger for Bing users. Either way, our results suggest that mandating choice screens on all devices (including those currently subject to a Bing default) may increase Google's overall market share.

Switch Bonus Group. We now analyze the market shares of baseline Google users in the Switch Bonus group. We first focus on participants who experienced Bing without Ranking Degradation or Ad Blocking. During the two-week incentive period, the Bing market share in this treatment group was 64 percent, as can be seen in Figure 3.

Strikingly, many of these users actively chose to continue using Bing after the end of the incentive period. During the week after Survey 2, the Bing market share among those who accepted our offer and qualified for the incentive payment was 38 percent, and it was still 35 percent during the last week of the experiment. This marks a stark departure from the Active Choice group. The Switch Group's resulting market share of 22 percent in the last week of the experiment (which includes participants who declined our offer) is substantially higher than the 2.5 percent in the Active Choice group. The only difference between these two groups is that Switch Group participants were exposed to Bing for two weeks. Our results hence suggest that their perceptions about Bing improved after exposure.

Ratings and exit survey responses help us interpret this result, and confirm that participants increased their relative preference for Bing. Table 4 shows the initial ratings in Survey 1 and the relative ratings change from Survey 1 to Survey 2 for the \$10 Switch Group (we also give an overview of responses in Appendix Figure A4).[10] A higher relative rating indicates a stronger preference for Bing. On average, Google users' ratings of Bing significantly improved in all dimensions except the number of ads shown. These changes are sizable if we compare them to the baseline survey ratings and their standard deviation. For instance, the effect on the overall quality rating between Google and Bing corresponds to a third of the initial gap between Google and Bing and more than half a standard deviation. In our exit survey, when we asked former Google users to give one or more reasons why they decided to stay with Bing after the incentive ended, 64 percent responded that Bing is better than expected, 59 percent that they got accustomed to it, 5 percent that they always wanted to use it, and 28 percent gave other reasons in a free form text response. The two most frequent responses (the first and second responses above) are consistent with our model, although they imply different interpretations of the quality term ζ. In the first case, ζ changes because people learned about Bing, and in the second because the utility of using Bing increased after getting accustomed to it. Our

[10] This table focuses on S10CC users only because we will later compare them to results that include the Ranking Degradation and Ad Blocking treaments (which only took place among S10 users).

Table 4: **$10 Switch Group: Change in Ratings for Bing Relative to Google**

	(1) Overall quality	(2) Relevant links	(3) Result page features	(4) Relevant ads	(5) AI chat	(6) Privacy	(7) Rewards	(8) Number of ads
Panel A: Baseline Google Users								
Δ Rating	0.403***	0.231***	0.511***	0.452***	0.489***	0.425***	0.353***	-0.077
	(0.074)	(0.078)	(0.090)	(0.081)	(0.077)	(0.063)	(0.066)	(0.058)
Rating at baseline: Mean	-1.475	-1.281	-1.158	-0.851	-0.421	-0.633	0.231	0.683
Rating at baseline: Sd	0.784	0.849	0.883	0.831	1.078	0.923	1.081	0.719
N	221	221	221	221	221	221	221	221
Panel B: Baseline Bing Users								
Δ Rating	-0.394*	-0.121	-0.303	-0.121	-0.303*	-0.121	-0.152	-0.091
	(0.213)	(0.212)	(0.211)	(0.161)	(0.171)	(0.155)	(0.152)	(0.118)
Rating at baseline: Mean	0.364	0.121	0.333	0.091	1.030	0.242	1.576	0.848
Rating at baseline: Sd	1.295	1.244	1.021	0.843	0.883	0.663	1.001	0.712
N	33	33	33	33	33	33	33	33

Notes: This table presents the average change (from Survey 1 to Survey 2) in ratings of Bing relative to Google on the search engine rating questions. The sample is participants in the $10 Switch Bonus group that passed the attention check, including those that did not accept the offer to switch. The survey questions in column 1 and columns 2–7, respectively, are "Overall, how would you rate the quality of Google relative to Bing?" and "How would you rate the quality of Google relative to Bing on the following dimensions?" Response options were "Bing is a lot better," "Bing is a little better," "They are about the same," "Google is a little better," and "Google is a lot better," coded as 2, 1, 0, -1, and -2, respectively. The survey question in column 8 is "How do you feel about the number of ads on Bing?" Response options were "way too many," "too many," "right amount," "too few," and "way too few," coded as 2, 1, 0, -1, and -2, respectively.*, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively. Standard errors clustered at the participant level.

counterfactual results are invariant to these two interpretations.[11]

Among baseline Bing users, the Google market share is 59 percent during the two-week incentivized period. We also find that an important fraction of these users decide to stay with Google: the Google market share among S10CC was 16 percent during the week after Survey 2, and it was still 15 percent during the last week of the experiment. However, unlike what we observe among baseline Google users, this market share is not statistically different from the 18 percent market share among Active Choice users. This suggests that Bing users are well-informed about Google's quality. The results from the rating survey are mostly consistent with this interpretation (Table 4): Bing users in general update less towards Google, and only the overall quality and AI chat updates are marginally significant.

Default Change. In the Default Change group, Bing's market share during the two-day incentive period was 73 percent, trending down gradually over the next few weeks: it was 50 percent seven days after Survey 1, and it was 44 percent on day 14.[12] This gradual decline is consistent with our stochastic model of inattention; in the absence of inattention, we would have expected to see a sudden drop in usage after the end of the incentivized period. The market share eventually stabilizes at 40 percent after four weeks, significantly higher than the 17 percent who still use Bing in the Switch Bonus group at that time.

There are multiple explanations for this persistently high Bing market share after the default change. First, some users may be persistently attentive, as suggested by the fact that the Bing market share in the Default Change group at the end of the experiment is significantly higher than the 17 percent in the Switch Bonus group. Second, users may have changed their perceptions of Bing's quality during the incentive period, consistent with the significant amount of learning we observe in the Switch Bonus Group. The results from our exit survey suggest that both explanations play a role: while 44 percent of participants reply that they still use Bing because they forgot to switch back or that switching back was too much effort, 35 percent of participants reply that they kept using Bing because they prefer it. The significant share of people who revise their perceptions about Bing after the default change suggests another important effect of Google's default agreements, which is to prevent users from learning about about the search engine that would otherwise be the default.

Price Responsiveness. Panel (b) of Figure 3 shows the market shares for the different price levels in the Switch Bonus groups. We see evidence of a clear price response during the incentive period. Even a \$1 payment leads to a 32 percent Bing market share—30 percent higher than the Active Choice market share—suggesting that an important fraction of users have weak preferences. A \$10 payment increases Bing's market share to 64 percent, and a \$25 payment results in 74 percent. The small additional gains in Bing's market share for higher prices imply that the distribution of willingness to accept is fat-tailed, with many

[11] One third possibility is a model where the utility of Google decreases after exposure to Bing as users get unaccustomed to Google. In counterfactuals, this model would result in the same market shares as our current model, but the effects of exposure to Bing on consumer surplus would be smaller.

[12] We do not observe any meaningful change in the time trend when these users received the placebo Survey 2.

people weakly and some people strongly preferring Google. Our stated preferences survey confirms this: 1.3 percent of users say that Bing is a lot better than Google, while 61 percent of users say that Google is a lot better than Bing (Appendix Figure A6 shows the full bar chart). Our model will accommodate this fat-tail pattern through a flexible functional form for the idiosyncratic preferences that allows for, but does not impose, such asymmetry.

There is a drop in all Bing market shares once the incentive period ends during Survey 2. Post-Survey 2 market shares are higher for higher payments, consistent with our interpretation that users exposed to Bing positively update their beliefs about its quality: in higher payment groups, a larger fraction of users update their beliefs, resulting in higher market shares after the incentive period ends.

4.4 Effects of Ranking Degradation and Ad Blocking

We now present the effects of Ranking Degradation and Ad Blocking on click-through rates (CTR), search engine ratings, and market shares. The analysis in this section is limited to baseline Google users in the $10 Switch Bonus group, since they are the only users that could have been randomized into Ad Blocking or Ranking Degradation.

Survey ratings. We first investigate the effects of these interventions on the search engine rating questions. Define Y_i as the difference in responses to a survey question on Survey 2 versus Survey 1, where more positive values indicate a more positive update about Bing, and let w_i^R and w_i^A be indicators for the Ranking Degradation and Ad Blocking groups, respectively. We estimate the following regression in the sample of participants in the $10 Switch Bonus group who passed the attention check:

$$Y_i = \tau^R w_i^R + \tau^A w_i^A + \mu + \varepsilon_i, \tag{9}$$

where τ^R and τ^A are the coefficients of interest, and μ is a constant.

Panel A of Table 5 presents results. The coefficients τ^R and τ^A capture the effects of the Ranking Degradation and Ad Removal treatments, respectively. Ranking Degradation significantly reduces the positive updating about Bing on overall quality, the relevance of links, the result page features, and the relevance of ads. Interestingly, Ranking Degradation's effect on the relevance of results is similar to (though a little larger than) the learning effect reported in Table 4. This indicates that Google users in Survey 1 expect the results of Bing to be (almost) as bad as we make them with Ranking Degradation. Ad Blocking significantly reduces the positive update about the relevance of ads and the result page features.

Click-through rates. We next investigate the effect of these interventions on clicks on the Bing results pages. To do this, we re-estimate equation (9) defining Y_i as various click-through rates. We limit the sample to users that accepted the offer on Survey 1 to switch to Bing in exchange for $10—otherwise they cannot click on Bing results.

Table 5: **Effects of Ranking Degradation and Ad Blocking on Bing Clicks and Quality Ratings**

	Panel A: $10 Switch Group: Change in Ratings of Bing Relative to Google							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var:	Overall quality	Relevant links	Result page features	Relevant ads	AI chat	Privacy	Rewards	Number of ads
Ad Blocking	-0.038	-0.050	-0.152*	-0.169**	-0.087	-0.071	-0.118*	0.090
	(0.070)	(0.072)	(0.086)	(0.072)	(0.076)	(0.061)	(0.065)	(0.060)
Ranking Degradation	-0.287***	-0.311***	-0.328***	-0.165**	-0.037	-0.099	0.065	-0.019
	(0.070)	(0.072)	(0.086)	(0.073)	(0.076)	(0.061)	(0.065)	(0.060)
Constant	0.425***	0.241***	0.583***	0.432***	0.479***	0.425***	0.365***	-0.050
	(0.063)	(0.066)	(0.077)	(0.068)	(0.067)	(0.054)	(0.057)	(0.051)
R^2	0.019	0.021	0.020	0.012	0.002	0.005	0.005	0.003
N	895	895	895	895	895	895	895	895

	Panel B: Clicks On Bing				
	(1)	(2)	(3)	(4)	(5)
Dep. var:	Original search rank	Ad clicks	CTR all	CTR organic	CTR organic (top)
Ad Blocking	0.033	-2.996***	-0.013	0.012	-0.017*
	(0.096)	(0.367)	(0.027)	(0.026)	(0.010)
Ranking Degradation	3.422***	1.867***	-0.057**	-0.075***	-0.153***
	(0.090)	(0.354)	(0.027)	(0.027)	(0.010)
Constant	1.640***	2.082***	0.361***	0.343***	0.237***
	(0.054)	(0.229)	(0.017)	(0.017)	(0.008)
R^2	0.666	0.123	0.008	0.014	0.294
N	646	646	636	636	636

Notes: This table presents estimates of equation (9), showing the effects of Ranking Degradation and Ad Blocking on ratings of Bing relative to Google on the search engine rating questions (in panel A) and Bing click outcomes (in panel B). In panel A, the sample is participants in the $10 Switch Bonus group that passed the attention check, including those that did not accept the offer to switch. The survey questions in column 1 and columns 2–7, respectively, are "Overall, how would you rate the quality of Google relative to Bing?" and "How would you rate the quality of Google relative to Bing on the following dimensions?" Response options were "Bing is a lot better," "Bing is a little better," "They are about the same," "Google is a little better," and "Google is a lot better," coded as 2, 1, 0, -1, and -2, respectively. The survey question in column 8 is "How do you feel about the number of ads on Bing?" Response options were "way too many," "too many," "right amount," "too few," and "way too few," coded as 2, 1, 0, -1, and -2, respectively. In panel B, the sample is further limited to users who accepted the offer to switch to Bing on Survey 1. The outcomes are the average original rank (before ranking degradation) of the Bing organic results the user clicked on (column 1), the count of ad clicks on Bing (column 2), Bing click-through rate (CTR) including both ad clicks and search link clicks (column 3), Bing CTR including only search link clicks (column 4), and Bing CTR including only the clicks associated with the first-ranked Bing search result (column 5). *, **, ***: statistically significant with 90, 95, and 99 percent confidence, respectively. Standard errors clustered at the participant level.

Panel B of Table 5 presents results. The first two specifications are manipulation checks that show that our Ranking Degradation and Ad Blocking treatments are working as expected. Columns 3–5 show the effects of Ranking Degradation on various click-through rates. These click-through rates reflect whether participants find the results useful and will also serve as our main measure of result relevance in the returns-to-scale analysis. Our results suggest that participants find the organic results after Ranking Degradation less useful: the overall click-through rate on organic links drops by 7.51 percent (Column 4), and the top link click-through rate falls by 15.3 percent (Column 5). Instead, participants are more likely to click on ads, which attenuates the overall effect on clicks (Column 3).

Ad Blocking, on the other hand, does not significantly affect participants' interaction with organic search results, nor does it much affect overall click-through rates.

Market shares. Figure 5 shows the effects of the Ranking Degradation and Ad Blocking treatments on Bing market shares. Let Y_{it} be Bing's market share, with t indexing periods specifically defined for this regression (and pre-registered as part of our pre-analysis plan): 0–14 days before Survey 1, 0–7 days after Survey 1, 8 days after Survey 1 to the day before Survey 2, 0–7 days after Survey 2, and >7 days after Survey 2. We estimate the following regression:

$$Y_{it} = \tau_t^R w_i^R + \tau_t^A w_i^A + \mu_t + \varepsilon_{it}. \tag{10}$$

We let the coefficients τ_t^R and τ_t^A vary by period t. We use robust standard errors, clustered by participant. Figure 5 shows the τ_t^R and τ_t^A coefficients for each period. We would expect the interventions to have zero effect before Survey 1 (when they are not yet active), a limited effect between Survey 1 and Survey 2 (when they are active but users are incentivized to stick with their choice) and a more stark effect after Survey 2 (on which all users make an unincentivized active choice after experiencing Bing with the respective modification). Both interventions reduce Bing's market share after the active choice in Survey 2. These effects, however, are not significant, despite the fact that the extension intervention changes participants' behavior on the result page as well as their relative rating of Bing versus Google. This result suggest that, when making choices, participants simply do not place enough weight on these attributes relative to other considerations, such as the interface and Bing rewards.

5 Model Estimation

We now explain how we use our experiment to estimate the model. First, we describe how we map the data to the model. Next, we detail how the experimental treatments identify the model parameters and outline the generalized method of moments (GMM) procedure we use for estimation.

We let t index (approximately) two-week periods. Period $t = 0$ is the 14 days before Survey 1, $t = 1$ is the 14 days between Survey 1 and when Survey 2 is sent, and $t = 2, 3, \ldots$ are successive 14-day periods after Survey 2.

Figure 5: **Effects of Ranking Degradation and Ad Blocking on Market Share**

(a) **Effects of Ranking Degradation**

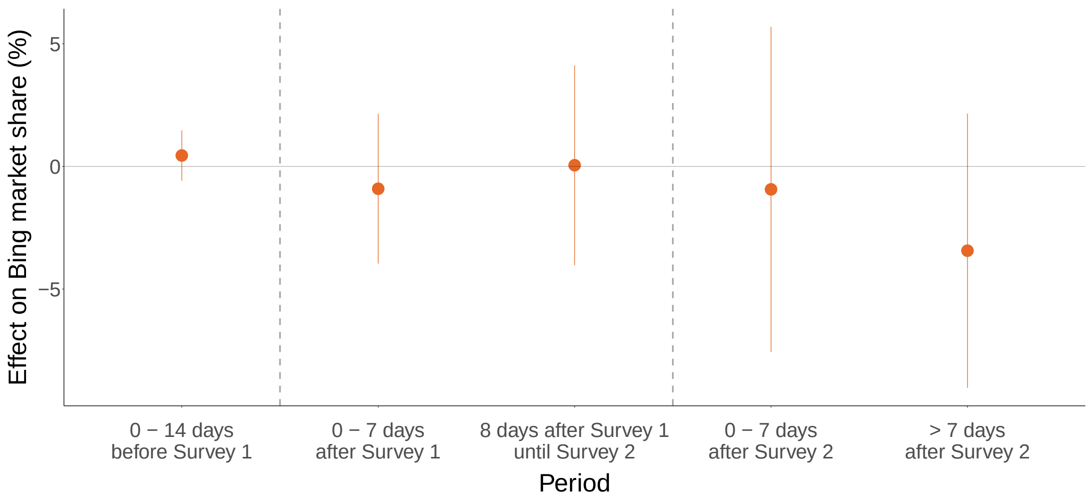

(b) **Effects of Ad Blocking**

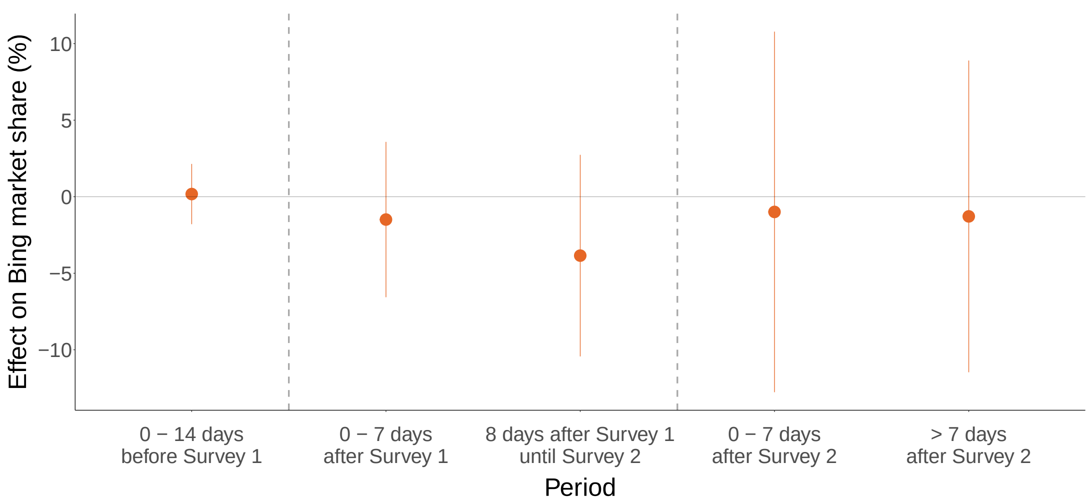

Notes: Panels (a) and (b) present the effects of the Ranking Degradation and Ad Blocking treatments on Bing market share by period of the experiment, estimated using equation (10). The sample includes only baseline Google users offered $10 to switch to Bing for the 14 days between Survey 1 and Survey 2. Whiskers indicate 95 percent confidence intervals. The dashed vertical lines mark the dates of the two surveys.

Figure 6: **Relation Between Experimental Market Shares and Parameters**

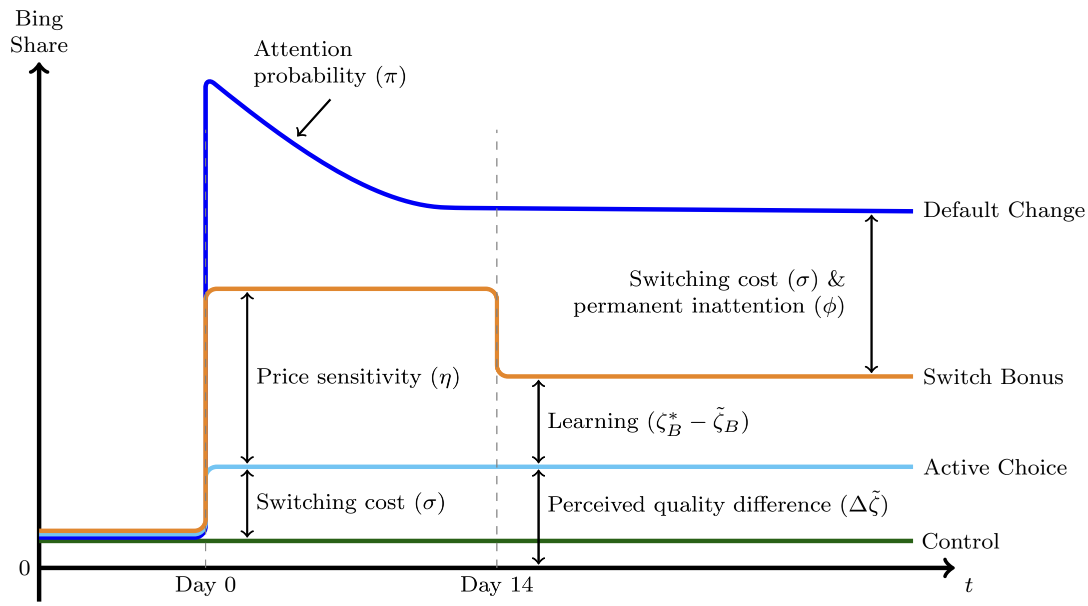

Notes: This figure illustrates, for Chrome users, the Bing market shares over time in each treatment condition. It also shows which elements identify each of the model parameters.

In the model, consumers make a binary choice between Bing and Google. In the data, a small number of people multi-home or occasionally use other search engines. For estimation, we say those users chose the search engine where they conducted the most searches.

Our model predicts browser-specific market shares, but our experimental treatment assignments depend on baseline search engine. For instance, baseline Bing users are paid to switch to Google in the Switch Bonus condition, and baseline Google users are paid to switch to Bing. To simplify the computation of market shares, we make a natural assumption that people who had already switched away from their browser's default search engine at baseline (e.g., Chrome users who search on Bing) would also do so if paid an incentive. This allows us to compute market shares at the browser level, as in our model.

5.1 Identification

We now provide an intuitive discussion of identification, using Figure 6, which provides a stylized representation of our data patterns for Chrome users. The arguments trivially extend to Edge users. We present formal identification arguments in Appendix C.1.

The perceived quality difference $\Delta\tilde{\zeta}$ is identified by the market share in the Active Choice group. Since these users must make a choice without any defaults or switching costs influencing them, a larger Bing share indicates higher perceived quality of Bing relative to Google.

The price sensitivity η is identified by the difference in market shares between the Switch Bonus group and the Active Choice group during the incentive period. Both groups make an active choice in Survey 1, but only the Switch Bonus group is paid.

Learning $\zeta_B^* - \tilde{\zeta}_B$ is identified by the difference in market shares between the Switch Bonus group and the Active Choice group after the incentive expires on day 14. When this incentive expires, users in the Switch Bonus group are asked to make another active choice. Hence, any difference in market shares between these two groups after day 14 must be driven by the Switch Bonus group's additional 14 days of experience with Bing.

The attention probability π is identified by the decaying pattern in the Bing market share of the Default Change group: the higher the attention probability, the more rapidly this market share will decay towards its long-run level.

The switching cost σ and the share of permanently inattentive users ϕ are jointly identified by two market share differences: the long-run difference between the Default Change group and the Switch Bonus group, and the difference between the Active Choice group and the Control group. Both differences are larger when these two sources of inertia are important. If there is no inertia, we would expect users in the Control group to already be with their preferred search engine (which they select in the Active Choice group). Similarly, if there is no inertia, whether or not there was a Default Change in the past should not affect a user's eventual choice of search engine.

While formally both σ and ϕ affect both differences, in practice the difference between Active Choice and Control is almost entirely driven by the switching cost σ. Intuitively, permanent inattention ϕ only affects users who would like to switch. Given the small market share of Bing in the Active Choice group (2.6%) , we can infer that few Control group users would like to switch. By contrast, users in the Default Group only switched to Bing because of the incentive payment, and therefore many of them would likely want to switch back. Hence, permanent inattention has a larger effect on the difference between Default Change and Switch Bonus than on the difference between Active Choice and Control. By contrast, the switching cost σ affects both differences roughly symmetrically as long as the density of marginal users is relatively similar in both Default Change and Active Choice.

The distribution of idiosyncratic preferences $\Delta\chi_i$ is identified from the market shares across the various Switch Bonus group payments and the Active Choice group: as we increase the incentive payment, we observe what fraction of users is willing to switch to Bing, and hence measure various quantiles of the idiosyncratic preference distribution. If we observed sufficiently many prices, we could identify the distribution of idiosyncratic preferences $\Delta\chi_i$ non-parametrically; given that our experiment includes only four different prices, our estimation below makes a parametric assumption.

5.2 Estimation

We estimate our model parameters by the generalized method of moments (GMM). The parameters that we estimate are the perceived quality difference $\Delta\tilde{\zeta}$, learning $\zeta_{-d}^* - \tilde{\zeta}_{-d}$, the price response η, the amortized switching cost $\sigma(1-\delta)$, the fraction of inattentive consumers ϕ, and the per-period probability of paying

Figure 7: **Identification of Idiosyncratic Preference Distribution**

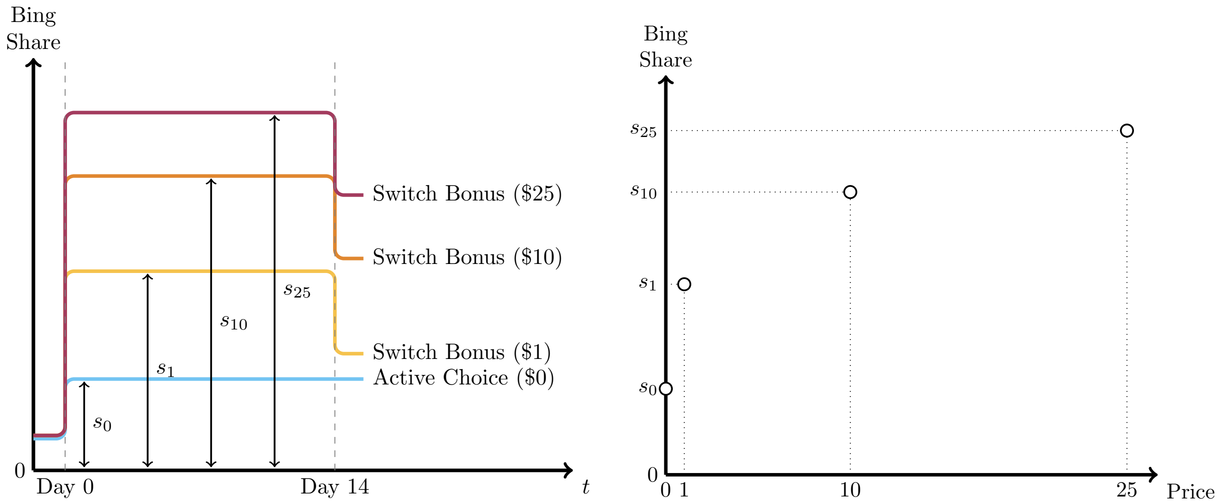

Notes: This figure illustrates, for Chrome users, the Bing market shares over time in each treatment condition (left). It also illustrates what fraction of users are willing to switch to Bing at any given price (right).

attention among attentive consumers π.

To estimate the distribution of idiosyncratic preferences, we assume it follows a shifted lognormal distribution. This parametric assumption has an important advantage relative to more commonly used distributions, such as normal or logistic. Depending on the parameters, the distribution can range from a symmetric distribution (approaching a normal distribution in the limit) to a heavily skewed distribution with a fat tail, like the one suggested by many users being unwilling to switch to Bing for a payment of $25. After normalizing the distribution, it has one remaining parameter $\gamma \in (0, \infty)$ that captures its skewness: $\gamma = 0$ represents a symmetric distribution, and the distribution becomes skewed as $\gamma \to \infty$, with a fat tail of customers that have a strong preference for Google.[13] Given our normalization, the preference shifters $\Delta\tilde{\zeta}$ and $\Delta\zeta^*$ represent differences from the value that would result in a market share of one half.

We refer to the full vector of parameters as θ. Theoretically, we can identify each element of θ separately for Edge and Chrome users. In practice, our small sample of Edge users forces us to pool estimation of most parameters to economize on power. Hence, we only allow the quality difference $\Delta\tilde{\zeta}$ and learning $\zeta^*_{-d} - \tilde{\zeta}_{-d}$ to differ between Chrome and Edge users. Since Figure 4 and Table 4 show that Bing users do not learn much about Google, and since the relevant market shares are noisy, we directly set $\zeta^*_G = \tilde{\zeta}_G$. We set the rest of the parameters to be the same for all users, no matter whether they are on Edge or on Chrome.

The moments that we use closely follow the identification arguments from Section 5.1. All moments are based on market shares that we observe in the data, which we present in Table 6. We now present a high-level description of the moments that we use, relegating further details to Appendix C.2.

[13] We normalize the distribution so that $S(-1) = 0$ and $S(0) = 1/2$. The PDF of the normalized distribution of $\Delta\chi$ is $\exp\left(-\log(x+1)^2/(2a^2)\right) / \left(\sqrt{2\pi}\gamma(1+x)\right)$.

Table 6: **Empirical Moments for Demand Estimation**

Description	(1) Formula	(2) Estimate	(3) SE
Baseline Chrome users			
Baseline and Control, $t \geq 0$	$\hat{s}^{C}_{-d,t\geq0}$	0.013	0.003
Active Choice group Bing share, $t \geq 1$	$\hat{s}^{A}_{-d,t\geq1}$	0.026	0.013
Default Change group Bing share, $t = *$	$\hat{s}^{D}_{-d,t=*}$	0.771	0.040
Default Change group Bing share, $t = 1$	$\hat{s}^{D}_{-d,t=1}$	0.470	0.047
Default Change group Bing share, $t = 2$	$\hat{s}^{D}_{-d,t=2}$	0.397	0.047
\$1 Switch Bonus group Bing share, $t = 1$	$\hat{s}^{S1}_{-d,t=1}$	0.337	0.056
\$25 Switch Bonus group Bing share, $t = 1$	$\hat{s}^{S25}_{-d,t=1}$	0.750	0.046
\$10 Switch Bonus (CC) group Bing share, $t = 1$	$\hat{s}^{S10}_{-d,t=1}$	0.677	0.032
\$10 Switch Bonus (CC) group Bing share, $t \geq 2$	$\hat{s}^{S10}_{-d,t\geq2}$	0.206	0.027
\$10 Switch Bonus (CR) group Bing share, $t \geq 2$	$s^{R}_{-d,t\geq2}$	0.158	0.024
\$10 Switch Bonus (BC) group Bing share, $t \geq 2$	$s^{B}_{-d,t\geq2}$	0.144	0.040
Baseline Edge users			
Baseline Google share, $t = 0$	$\hat{s}_{-d,t=0}$	0.273	0.104
Active Choice group Google share, $t \geq 1$	$\hat{s}^{A}_{-d,t\geq1}$	0.368	0.112
\$10 Switch Bonus (CC) group Google Share, $t = 1$	$\hat{s}^{S10}_{-d,t=1}$	0.709	0.106
\$10 Switch Bonus (CC) group Google Share, $t \geq 2$	$\hat{s}^{S10}_{-d,t\geq2}$	0.389	0.113

Notes: This table presents the empirical moments used for the demand estimation procedure described in Section 5. Standard errors clustered at the participant level.

The first set of moments are simply market shares: the baseline market share s_{-d0}, the Active Choice market share $s^{A}_{-d,t\geq1}$, the market shares for the Switch Bonus group during the incentivized period at different prices $s^{S1}_{-d,t=1}$, $s^{S10CC}_{-d,t=1}$, and $s^{S25}_{-d,t=1}$, and the post-Survey 2 market shares of the Switch Bonus group under different interventions $s^{S10CC}_{-d,t\geq2}$, $s^{S10RC}_{-d,t\geq2}$, $s^{S10CA}_{-d,t\geq2}$, and $s^{S10RA}_{-d,t\geq2}$. Following the arguments in Section 5.1, these shares identify the distribution of idiosyncratic preferences, perceived differences in quality, learning, the price response, and quality preferences.

To identify the attention probability π, we exploit the market shares of the Default Change group right after Survey 1, after one week, and after two weeks. Based on these three shares, we write out a moment condition that captures how quickly the market share converges to its long-run value (see Appendix C.2). Finally, to identify switching costs and inattention, we need moments for $s^{A}_{-d,t\geq0} - s^{C}_{-d,t\geq0}$ and $s^{D}_{-d\infty} - s^{S10CC}_{-d,t\geq2}$. We already included moments corresponding to s_{-d0} (which is the same as $s^{C}_{-d,t\geq0}$), $s^{A}_{-d,t\geq1}$, and $s^{S10CC}_{-d,t\geq2}$, so we include an additional moment that characterizes $s^{D}_{-d\infty}$ as a function of shares that we observe in the data (see Appendix C.2).

Our estimates are given by

Table 7: **Demand Parameter Estimates**

Description	(1) Formula	(2) Estimate	(3) SE
All users			
Distribution shape	a	2.918	0.546
Price response	η	0.326	0.086
Permanent inattention	ϕ	0.336	0.058
Attention probability	π	0.826	0.150
Amortized switching cost	$\sigma(1-\delta)/\eta$	\$0.004	0.007
Baseline Chrome users			
Bing preference shifter	$\Delta\tilde{\zeta}/\eta$	-\$3.058	0.796
Learning	$(\zeta^*_{-d}-\tilde{\zeta}_{-d})/\eta$	\$0.258	0.180
Ad load response	$\alpha(a^{CA}_{-d}-a^{CC}_{-d})/\eta$	-\$0.127	0.122
Quality response	$\rho(r^{RC}_{-d}-r^{CC}_{-d})/\eta$	-\$0.100	0.101
Baseline Edge users			
Google preference shifter	$\Delta\tilde{\zeta}/\eta$	-\$7.673	0.679
Learning	$(\zeta^*_{-d}-\tilde{\zeta}_{-d})/\eta$	\$0	-

Notes: This table presents the parameter estimates from the demand estimation procedure described in Section 5. Standard errors are clustered at the participant level.

$$\hat{\theta} = \text{argmin}_{\theta} G(\theta)'\Omega G(\theta), \tag{11}$$

where m indexes different moments, $G_m(\theta)$ is the average of $g_{im}(\theta)$ across users, and Ω is a weighting matrix. We use two-step GMM, where we first set $\Omega = Cov(G(\theta))^{-1}$ for some arbitrary θ, and then, in our second step, we set Ω to be the optimal weighting matrix given the initial set of estimates.

5.3 Demand Parameter Estimates

Table 7 presents the demand parameter estimates. We find that the distribution of idiosyncratic preferences is very skewed (i.e., a high value of a). As we can see in Figure 8, such a skewed distribution is entirely consistent with (and identified by) the price response patterns that we described in Section 4.3: there is a large density of users that are close to indifferent between both search engines, essentially no users with strong preferences in favor of Bing, and a fat tail of users with strong preferences in favor of Google.

Such a skewed distribution has two important implications. First, a given change in price or quality has a large impact on market shares when the marginal users are those that are close to indifferent between both search engines (when the market share of Google is high), but it only has a modest impact when

Figure 8: **Distribution of Payment Required to Switch to Bing**

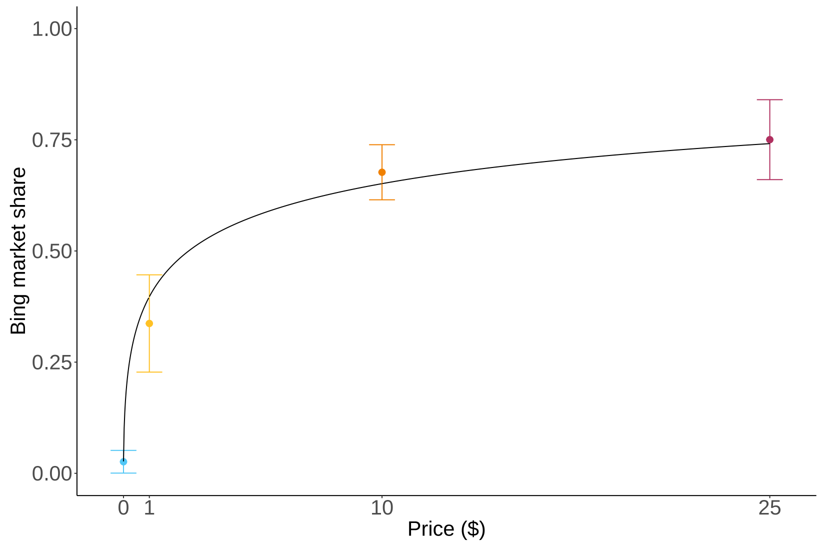

Notes: This figure shows our estimates of the fraction of users willing to switch their default search engine to Bing for two weeks (on the vertical axis) for a given payment (on the horizontal axis.) We exhibit means with associated standard errors (dots and associated error bars) and the distribution that we fit to these market share moments (the solid line). The plot confirms that our parametric assumption fits the data well.

the marginal users belong to the fat tail of users with strong preferences in favor of Google (when the market share of Google is low). This feature drives many of our findings below. Second, the combination of a fat tail and permanent inattention implies large consumer surplus losses due to agents with strong preferences that nevertheless use an undesired search engine because they are inattentive. For that reason, when computing consumer surplus, we censor idiosyncratic preferences at $25 to avoid our results being driven by extrapolation far beyond the price offers that S group participants received.

We now discuss parameters for Chrome users as well as the rest of the parameters for all users—which are mainly identified from data on Chrome users because they account for most participants in our sample. We use our estimate of the price response η to interpret all other parameters in units of dollars per two-week period. For Chrome users who have not used Bing, we observe a negative Bing preference shifter $\Delta\tilde{\zeta}$ (equivalent to a payment of $3.06 per two-week period), consistent with the low Bing market share in the data. After exposure to Bing, its perceived utility increases by $0.26 per two-week period. Updating perceptions about Bing affects those users who are close to indifferent between both search engines, where there is a large density of users. Thus, as we note above, a small change in the perception of quality is enough to generate a large change in market shares.

The estimated switching costs $\sigma(1-\delta)$ are small (0.4 cents), given the small effect on market shares of the Active Choice intervention—which, as we explain in Section 5.1, are mainly driven by switching costs rather than inattention. On the other hand, we find that inattention plays an important role as 34 percent of users are permanently inattentive. The per-period attention probability π of 83 percent means that users who are not permanently inattentive make attentive choices frequently. Our estimates for the preferences for quality indicate that Ad Blocking decreased utility as much as a price decrease of \$0.13, and Ranking Degradation caused a utility decrease of \$0.10. Neither of these two parameters is statistically significant.

For Edge users, we find a negative Google preference shifter $\Delta\tilde{\zeta}$ that is equivalent to a payment of \$7.67 per two-week period, consistent with the high share of Bing among Edge users.

6 Economies of Scale in Data

In this section, we are interested in estimating the causal effect of observing additional past searches of a given search term on the current relevance of the results presented to users who search for that term. To that end, we use data on search terms and clicks from Bing. As is standard industry practice, we will use click-related outcomes as our measure of the relevance of the links presented on a search results page. In particular, our measure of relevance is the probability that a user clicks on the top-ranked result. Henceforth, we will refer to search terms as queries and instances where a user enters a query as searches.

An important input to our model is the extent to which the relationship between data and search result relevance exhibits diminishing returns to scale. Google argues that the effect of additional search data on improving result rankings diminishes quickly, implying that increased scale has little effect on the relevance of results (e.g., Varian 2015). This argument is less convincing if there are many *tail* queries (i.e., queries with few searches) for which additional click and query data may still be valuable. Indeed, examining all searches made on Bing over 12 months in 2021 and 2022, we find that the distribution of queries exhibits a long tail: Figure 9 shows what fraction of searches accrues to queries with different occurrence rates. More than 38.7 percent of searches are for rare queries that are searched less than 100 times. For this reason and for reasons of identification that we explain below, we will focus on new queries.

For our analysis, we randomly sampled 43,991 new search terms, i.e., search terms for which Bing had no search record between January 2021 and January 2022. For these search terms, we recorded each search in the subsequent year (i.e., between January 2022 and January 2023).[14]

6.1 Empirical Strategy

We are interested in recovering the effect of the number of searches n_{qt} on result relevance r_{qt} for query q at time t:

$$r_{qt} = f(n_{qt};\theta) + \varepsilon_{qt}. \tag{12}$$

[14] The resulting panel dataset at the (query, search)-level records (i) the date of each search, (ii) various click-related outcome measures (was there any click, was the top result clicked, was there any click from which the user did not immediately return) for each search, and (iii) an identifier for the URL that was top-ranked for these searches. We report summary statistics for this dataset in Table A4 in Appendix D.1. Optimally, we would like to subset to queries which Bing has never seen, but we are limited by Bing's retention period of 24 months.

Figure 9: **Distribution of Views Across Queries**

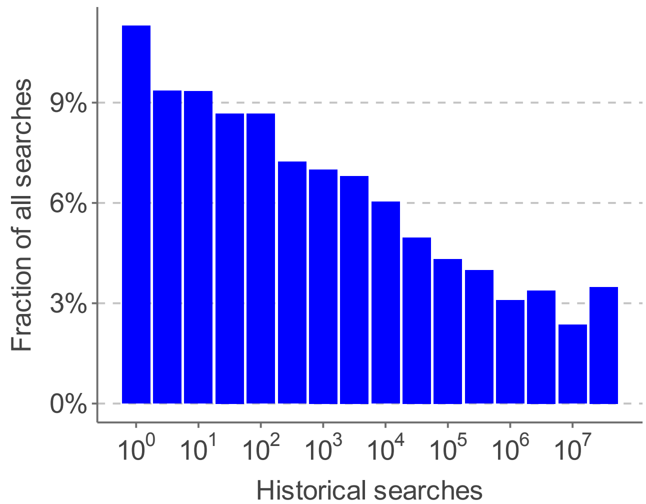

Notes: This figure provides the fraction of all searches for queries of a certain popularity. We count all searches made on Bing in the period between October 1st, 2021 and October 1st, 2022, and group queries by how often they were searched for. The shape of the distribution supports the commonly-held assumption that there are a large number of tail queries each associated with a lower number of searches but jointly responsible for a considerable share of searches on Bing.

In estimating this relationship, we need to address two potential confounders. First, there could be persistent differences in the difficulty of serving results across queries. For instance, navigational queries ("facebook.com") are both very common and easy to answer even for a search engine with only a small fraction of Google's data. Thus, a simple cross-sectional comparison may bias results in favor of finding returns to scale. To control for such persistent differences in the difficulty of serving results across queries, we use previously unseen queries over time as they gather their first searches and control for query fixed effects (as in He et al., 2017).

However, this strategy does not address time-varying confounders. Such confounders may arise due to compositional changes in the types of users over the lifetime of a query. For instance, suppose a query is concerned with a news event. For users who arrive early in the query's lifetime, the event may be more newsworthy, and they are hence more likely to click on results. By contrast, users that arrive later may already know of the event and just want to verify its date—something that can be accomplished without the need to click. If we did not account for time-varying confounders, this difference in user types over a query's lifetime could bias our estimation.

To deal with such time-varying confounders, we rely on a key identifying assumption, which seems plausible in this context: any *causal* effect of more search data on clicks has to be due to the results that the search engine serves and the order in which it presents them.[15] In Appendix D.3, we formalize an identifi-

[15]This assumption would be violated, for example, if the search engine also uses data to improve the interface of the search

Figure 10: **Event Study Illustrating Effect of Ranking Change on Click-Through Rates**

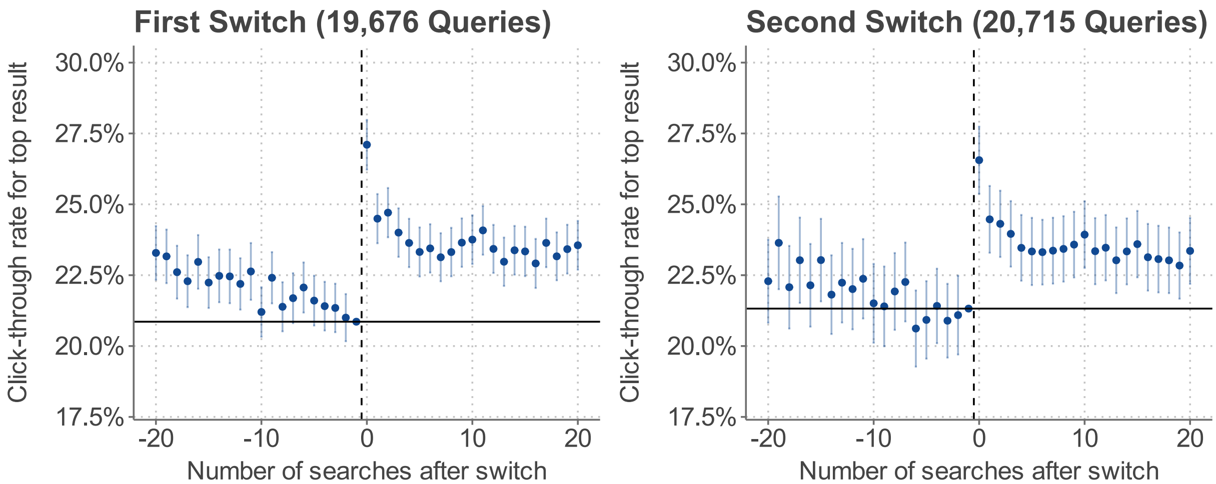

Notes: We regress the CTR on query fixed effects and a set of dummies measuring how many searches happened since the k-th switch in the top-ranked result. These figures report the these dummies, providing evidence that CTR increases at the time of a ranking change but otherwise follows a secular decline. The secular decline motivates our identification strategy (as it suggests the presence of a potential confounder), and the positive effect of ranking changes suggests that rankings do affect CTR.

cation strategy that builds on this assumption by exploiting the front door criterion (Pearl and Mackenzie, 2018; Imbens, 2020; Bellemare et al., 2024) to purge confounding variation in click-through rates. Here, we restrict ourselves to present the event studies that we leverage in this approach.

Figure 10 shows click-through rates against the number of searches a given query has had since the first or second time its top-ranked result changed. Updates to the ranking lead to an increase in click-through rates, which indicates that over time the search engine learns to serve more useful results. We also observe a general downward trend in click-through rates unrelated to the results the search engine serves. Intuitively, just like this event study, our identification strategy in the appendix isolates the variation in click-through rates that is driven by changes in the search result order.

6.2 Estimation

We now estimate the relationship between search result relevance (as proxied by the click-through rate) and the number of prior searches for a given query q on date t. For simplicity, we specify a functional form that allows for different relationships between data and CTR, including linear ($\theta = 0$), logarithmic ($\theta = 1$), and more concave than logarithmic ($\theta > 1$):

$$r_{qt} = \alpha_q + \beta \frac{(n_{qt})^{1-\theta} - 1}{1-\theta} + \varepsilon_{qt}. \tag{13}$$

results. However, to cause concern for our identification strategy, such interface adjustments would have to be query specific, which seems less plausible.

Table 8: **Economy of Scale Estimates**

Description	(1) Parameter	(2) Estimate	(3) SE
Click-through rate at inception	α	0.1811	-
Value of additional data	β	0.0056	(0.0007)
Shape of returns from data	θ	0.9458	(0.0318)

Notes: This table presents our parameter estimates for the relationship between the amount of data and the relevance of search results, as measured by the click-through rate. See equation (13) for the functional form specification.

As we show below, it fits the data quite well.

To ensure that we only use variation in the CTR that derives from changes in the organic ranking, the specification we run takes the form of equation 13 but the left hand side is $\hat{r}_{qt}$, a prediction of the CTR that is based only on the top result (see Appendix D.3). This process requires estimation of a non-linear model with high-dimensional fixed effects, for which we develop a simple procedure in Appendix D.4. Since our equilibrium model does not distinguish between different queries, we choose an overall intercept α by matching the average click-through rate predicted by our model to that in our experimental data.[16]

Table 8 presents the resulting estimates. We can strongly reject the null hypothesis that data does not matter ($\beta = 0$). As to the shape of the returns from data, the estimates strongly point towards a log-linear relationship between CTR and number of searches. Thus, doubling the amount of searches will lead to a fixed increase in CTR, no matter the starting point. Our estimates imply that if Bing had access to Google's data, CTR would increase from 23.5 percent to 24.8 percent.

We show the fit of the model in Figure 11, where we compare the model predictions to a binscatter plot of the data. The model fits well across many orders of magnitudes. Furthermore, our estimates are not too different from those in prior studies. [17]

We can use our fitted model to calculate counterfactual average click-through rates on Bing if its market share were to be multiplied by λ: the counterfactual click-through rate $\tilde{r}(\lambda)$ after such an increase in market share is given as a function of the status-quo click-through rate $\bar{r}$ by

$$\tilde{r}(\lambda) = \alpha - \frac{\beta}{1-\theta} + \lambda^{1-\theta} \cdot \Big(\bar{r} - \alpha + \frac{\beta}{1-\theta}\Big). \tag{14}$$

We use this expression in the computation of long-run counterfactuals below.

[16]This average is computed by integrating over the complete query frequency distribution and weighting each query by its number of searches. While we would optimally use lifetime query frequencies (i.e., query q has seen n views since inception), we actually observe query frequencies over a one year period.

[17]He et al. (2017, Fig. 4 and 5) find that going from 300 to 1,900 searches yields a CTR improvement of about 2 percentage points Our estimates imply an improvement of 1.5 percentage points. Schaefer and Sapi (2023, Fig 5(c)) find that going from 1 to 5,000 searches yields a CTR improvement of 3-5 percentage points. Our estimates imply a slightly larger improvement of 6.1 percentage points.

Figure 11: **Model Fit for Returns to Scale**

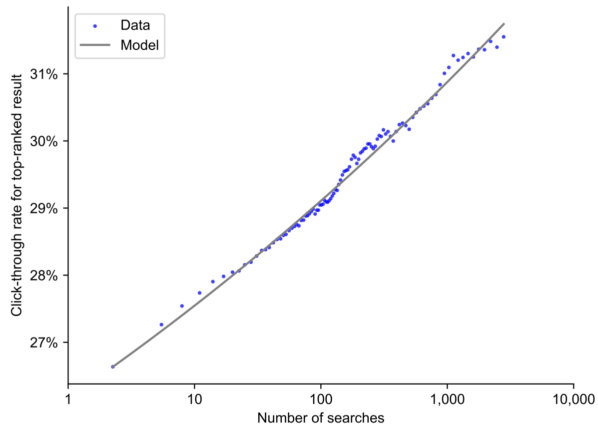

Notes: We present the model fit for equation (13) by exhibiting both the model-predicted CTR as a function of the number of prior searches for a query as well as a binscatter of the actually observed CTR as a function of the number of prior searches. The model estimates a linear relationship between CTR and the logarithm of the number of prior searches, and this relationship fits the observed data well.

Our approach to estimating economies of scale has several limitations. First, our model is estimated exclusively off new queries. This could be problematic if the effect of data differs between new queries and other types of queries. Second, extending our model to queries with many views requires extrapolation as new queries typically do not reach such high levels of popularity. Still, the fit of our model mitigates this concern. Lastly, the results above do not allow for cross-query learning. In Appendix D.6 we incorporate cross-query learning effects with quantitatively similar results.

7 Counterfactuals

To compute counterfactuals, we adjust the fraction of Chrome and Edge users from our experimental sample to reflect the observed US desktop market shares, which decreases Chrome's share and increases Edge's share.

Direct Effects We first analyze the direct impact of different scenarios on demand without accounting for returns to scale from data. Our results should be thought of as short-run effects, before the relevance of the results presented by both search engines changes in response to changes in market shares.

Table 9 shows our results. As a benchmark, we compare all counterfactuals to a *Status Quo* scenario, in which demand behaves just as in the control group, and market shares are given by $s_{-d} = (1-\phi)S\left(\Delta\tilde{\zeta} - \sigma(1-\delta)\right)$. We show two outcomes for each counterfactual: consumer surplus and market shares. We understand market shares not merely as a descriptive statistic but as an important proxy for

Table 9: **Direct Counterfactuals**

Panel A: Benchmarks						
	Combined		Chrome		Edge	
Description	(1) Google share (%)	(2) CS gain ($/user-year)	(3) Google share (%)	(4) CS gain ($/user-year)	(5) Google share (%)	(6) CS gain ($/user-year)
Status Quo	88.9	0.00	98.8	0.00	22.2	0.00
No Frictions	73.8	6.01	79.8	0.87	33.4	40.68
Active Choice	89.1	5.35	97.3	0.10	33.4	40.68
Correct Perceptions	78.4	0.46	86.7	0.53	22.2	0.00

Panel B: Policy Interventions						
	Combined		Chrome		Edge	
Description	(1) Google share (%)	(2) CS gain ($/user-year)	(3) Google share (%)	(4) CS gain ($/user-year)	(5) Google share (%)	(6) CS gain ($/user-year)
Choice Screen	87.6	0.09	97.3	0.10	22.2	0.00
Bing Default	48.9	-70.92	52.9	-81.44	22.2	0.00
Bing Default + Delayed Choice Screen	72.1	0.06	79.6	0.07	22.2	0.00
Bing Payments ($10)	51.3	108.78	56.4	93.76	17.0	209.99

Notes: This table presents counterfactual simulation results that only account for direct effects i.e., the effects before taking into account any potential change in result relevance implied by our economy of scale analysis. CS gain means consumer surplus gain, in $/consumer-year. Panel A presents hypothetical counterfactuals that serve as benchmarks. Panel B presents counterfactuals that represent policies.

the total welfare of all market participants. This is because market shares affect both competition between search engines for advertisers and search engine quality investment decisions. While we do not model these effects explicitly, their likely presence means that policies that can equalize market shares without significantly negatively affecting consumer surplus are desirable. Columns 1 and 2 present aggregate results for the whole market. Columns 3-6 then break down results into Chrome and Edge users.

To decompose the economic effect of different frictions, we first consider certain benchmark counterfactuals that cannot be attained with policies (Panel A). In our first scenario, *No Frictions*, we shut down all demand-side frictions: all users are attentive, they are informed about the true quality of search engines, and there is no switching cost. Therefore, market shares are given by $s_{-d} = S(\Delta\zeta^*)$ and only depend on the true quality of search engines. We find a sizable increase in the market share of Bing from 11 percent to 26 percent, though even in this frictionless world, Google's share would remain large. Consumer surplus increases by $6.01 per consumer per year. While unlikely in practice, this scenario is useful as a benchmark where the market achieves the social optimum while not yet accounting for scale economies.

Our next two counterfactuals decompose these effects into information frictions due to misperceptions of quality and due to sources of inertia—primarily permanent inattention, as switching costs are negligible. In the *Active Choice* scenario, we shut down switching costs and inattention. Therefore, demand behaves just as in the Active Choice treatment, and market shares are given by $s_{-d} = S(\Delta\tilde{\zeta})$. Rather than decreasing, Google's market share increases (although only by 0.2 percentage points): some Chrome users switch to Bing, but this is offset by a slightly larger number of Edge users switching to Google. Despite the small effect on shares, *Active Choice* does increase consumer surplus (by 89 percent of the gains from *No Frictions*) because more people use their preferred search engine. The main implication is that, although sources of inertia are an important friction—as evidenced by their impact on consumer surplus—they, by themselves, are not the reason why Google has such a high market share.

In the next counterfactual, *Correct Perceptions*, all users are perfectly informed about the true quality of both search engines. However, unlike the *No Frictions* counterfactual, users are still subject to switching costs and inattention. The market share is thus given by $s_{-d} = (1-\phi)S(\Delta\zeta^* - \sigma(1-\delta))$. Google's overall market share decreases by 10.5 percentage points. This decrease highlights that correcting consumers' misperceptions is a key element necessary for policy interventions to significantly reduce Google's market share.

Our remaining counterfactuals evaluate the effects of proposed policy interventions (Panel B). We first consider the effect of a *Choice Screen* like the one introduced by Android in Europe after the European Commission's 2018 decision. On Chrome, demand behaves as in our *Active Choice* benchmark. Edge user demand, on the other hand, remains unchanged from *Status Quo*. We can see that Google's market share decreases only by 1.2 percentage points, with only a slight increase in consumer surplus of \$0.09. These results show that choice screens have a limited effect on market shares, consistent with Decarolis et al. (2023), who documented minor effects following the introduction of choice screens in the EU.

Policies are more successful in moving market shares when they expose a larger fraction of the population to the alternative search engine, thus reducing misperceptions. In *Bing Default*, the default search engine upon installing a new browser is Bing on all browsers. This change approximates proposed remedies that would ban Google from bidding for default positions on Chrome, allowing Bing to outbid other competitors. Hence, while demand on Edge remains unchanged from *Status Quo*, both switching costs and persistent inattention now operate in favor of Bing for Chrome users. The counterfactual market share for Chrome users is given by $s_{-d} = \phi + (1-\phi)S(\Delta\zeta^* + \sigma(1-\delta))$. Google's overall market share declines by 40 percentage points. However, this intervention reduces consumer surplus by a significant \$70.92 per consumer per year because an important fraction of users are now defaulted into an undesirable search engine. The effect is so big because many users with strong preferences for Google use Bing due to permanent inattention.[18]

Our counterfactuals suggest that two of the most commonly proposed policies to curb Google's dominance have important drawbacks. Choice screens increase consumer surplus, but they barely move the

[18] As we explain in Section 5.3, we censor idiosyncratic preferences at \$25, so this value is not driven by the fat tail of preferences that we estimate. These numbers should thus be thought of as a conservative lower bound for what the actual welfare effects could be.

needle in terms of market shares. Default changes significantly impact market shares, but they come at the cost of a large reduction in consumer surplus. We now consider a policy that aims to achieve the best of both policies. In the *Bing Default + Delayed Choice Screen* counterfactual, Bing becomes the default search engine on all platforms (as in *Bing Default*), and users make an active choice after two weeks. The initial default allows users to learn about Bing, but the subsequent active choice avoids potentially large welfare losses for permanently inattentive consumers who prefer Google. During the first two weeks, market shares are the same as in *Bing Default*, but afterward, they shift to mimick the *No Frictions* counterfactual for Chrome users and the *Choice Screen* counterfactual for Edge users. To compute consumer surplus, we give weight $2/208$ to the consumer surplus before the choice screen and weight $206/208$ to the consumer surplus after the choice screen. These weights approximate what would happen if consumers reinstalled their browser every four years (the frequency with which people typically buy a new computer), with the default resetting to Bing upon reinstallation. This policy reduces Google's market share by 17 percentage points while barely affecting consumer surplus (which increases by \$0.06).[19] While this counterfactual might not be easily implementable, it highlights the key elements necessary for successful policy—exposing users to other search engines while preserving choice—and demonstrates the magnitude of potential gains.

Finally, we consider a *Bing Payments* counterfactual in which Bing pays users \$10 (on top of their existing Bing rewards program) to induce them to use Bing more.[20] This allows us to determine the extent to which payments can level the playing field in favor of Bing. Shares are given by $s_{-d} = (1 - \phi)S\left(\Delta\tilde{\zeta} - \sigma(1-\delta) + \eta p_B\right)$, where the payment $p_B = 10$ is positive for Chrome users but negative for Edge users. The payments achieve a large increase in Bing's market share. While they also achieve a large increase in consumer surplus, this is mainly because of the payment consumers receive.

Equilibrium effects We now analyze the effect of counterfactuals in equilibrium, accounting for economies of scale in data. We endogenize search result relevance using our model from Section 6 to capture the fact that search engines with a larger market share can use their improved access to user-generated data to improve the relevance of their results. Before conducting this analysis, we first define a market equilibrium.

Let the following function describe demand in counterfactual $\mathscr{C}$:

$$s = D(p, \zeta^*, \tilde{\zeta}; \mathscr{C}). \tag{15}$$

In this expression, s is the vector of market shares for both search engines, p is the vector of prices, ζ^* is the vector of true qualities, and $\tilde{\zeta}$ is the vector of the qualities perceived by agents who have not experienced search engines. This demand function differs across counterfactuals $\mathscr{C}$. For the *Status Quo,* for instance, market shares are given by $s_{-d} = (1-\phi)S(\Delta\tilde{\zeta} - \sigma(1-\delta))$, and they are $s_{-d} = S(\Delta\tilde{\zeta})$ for the *Active Choice*

[19] While the sign of the impact on consumer surplus is sensitive to the frequency at which we assume browser re-installation, the small magnitude of the effect is robust. Hence, we interpret this finding as suggesting one way for policy to affect market shares without significantly reducing consumer surplus.

[20] Such payments may be hard to implement in practice as people may be tempted to create multiple accounts and perform unnecessary searches to obtain larger payments.

counterfactual. Appendix E gives full expressions for all counterfactuals.

Based on our model of economies to scale from Section 6, the true quality is described by the following relationship:

$$\zeta^* = Z(s;\mathscr{C}). \tag{16}$$

In counterfactuals in which there is no data sharing, the function $Z(s;\mathscr{C})$ is given by the expression $\zeta_j^* = \alpha a_j + \rho \times \tilde{r}(s_j/\hat{s}_j, \bar{r}) + \xi_j$, where $\tilde{r}(\cdot)$ is as defined by equation 14 and $\bar{r}$ is the average click-through rate that we measure for Bing. This follows from substituting our results about click-through rates from Section 6 into our definition of quality $\zeta_j^* := \alpha a_j + \rho r_j + \xi_j$. In counterfactuals in which there is data sharing, each search engine's quality is as if they had a market share of one, so $Z(s;\mathscr{C})$ is as derived from the expression $\zeta_j^* = \alpha a_j + \rho \times \tilde{r}(1/\hat{s}_j) + \xi_j$.

An equilibrium consists of a joint solution of equations 15 and 16 in s and ζ^*. Let $s^{\text{eq}}(\mathscr{C})$ and $\zeta^{\text{eq}}(\mathscr{C})$ denote the equilibrium shares and qualities given counterfactual $\mathscr{C}$.[21] From these equilibrium quantities, we can compute consumer surplus in equilibrium from the demand function $D(p, \zeta^{\text{eq}}(\mathscr{C}), \tilde{\zeta}; \mathscr{C})$. Appendix E presents the exact expressions we use to compute counterfactuals and consumer surplus. Since our estimation of returns to scale in data relies on strong assumptions, this second set of counterfactuals should be interpreted as a speculative exercise that attempts to get a sense of the magnitude of the effects that may arise in equilibrium.

Table 10 presents results from our equilibrium counterfactuals. The first four rows of Panel A and the first five rows of Panel B analyze the same counterfactuals as in Table 9, while accounting for equilibrium effects. We find that economies of scale from data reinforce the direct effects of interventions on market shares by around one tenth. Additionally, we observe that economies of scale have only limited effects on consumer surplus: there are winners (Bing users) and losers (Google users), so the net effects are small. These limited effects can be explained by our estimates of the two forces driving network externalities, both of which are small: Section 5 finds moderate economies of scale, and Section 5.2 measures a weak demand response to quality.

We also consider additional counterfactuals in which regulators mandate Google to share data with its competitors, as proposed by certain antitrust authorities. This allows Bing to exploit the data from all users, so the relevance of the results presented by Bing is what it would be if Bing was a monopolist and could observe data from all users in the market. In the *Data Sharing* counterfactual (Panel B), we analyze what happens if regulators mandate data sharing without any additional intervention. The effects are almost negligible: Google's market share goes down by 0.002 percentage points, and consumer surplus increases by \$0.07. This policy is almost entirely ineffective because, although the additional data allows Bing to improve its quality, most users are unaware of the improvement and therefore stay with Google. Three

[21] In principle, there could be multiple equilibria: the economies of scale could be so big that the market either "tips" towards Google or Bing. In practice, that is not the case because we measure small economies and a limited demand response to search result relevance.

Table 10: **Equilibrium Counterfactuals**

Panel A: Benchmarks

	Combined		Chrome		Edge	
	(1)	(2)	(3)	(4)	(5)	(6)
Description	Google share (%)	CS gain ($/user-year)	Google share (%)	CS gain ($/user-year)	Google share (%)	CS gain ($/user-year)
Status Quo	88.9	0.00	98.8	0.00	22.2	0.00
No Frictions	73.5	6.02	79.4	0.86	33.4	40.80
Active Choice	89.1	5.35	97.3	0.10	33.4	40.67
Correct Perceptions	78.2	0.47	86.5	0.52	22.2	0.12
Correct Perceptions + Data Sharing	77.9	0.56	86.2	0.57	22.1	0.45
No Frictions + Data Sharing	73.1	6.12	79.0	0.93	33.4	41.05

Panel B: Policy Interventions

	Combined		Chrome		Edge	
	(1)	(2)	(3)	(4)	(5)	(6)
Description	Google share (%)	CS gain ($/user-year)	Google share (%)	CS gain ($/user-year)	Google share (%)	CS gain ($/user-year)
Choice Screen	87.6	0.09	97.3	0.10	22.2	0.02
Bing Default	48.5	-70.81	52.4	-81.36	22.2	0.27
Bing Default + Delayed Choice Screen	72.0	0.08	79.4	0.07	22.2	0.17
Bing Payments ($10)	51.3	108.88	56.4	93.83	17.0	210.28
Data Sharing	88.9	0.07	98.8	0.01	22.1	0.46
Data Sharing + Bing Default + Delayed Choice Screen	71.6	0.18	79.0	0.14	22.1	0.44

Notes: This table presents counterfactual simulation results that account for equilibrium effects: accounting for the changes in result relevance implied by our economy of scale analysis. The click-through-rate (CTR) used in the calculation of the equilibrium effects is the average consumer-level click-through-rates associated with top organic link clicks. Panel A presents hypothetical counterfactuals that serve as benchmarks. Panel B presents counterfactuals that represent policies.

counterfactuals consider alternative data sharing scenarios in which users have correct perceptions*: Correct Perceptions + Data Sharing* and *No Frictions + Data Sharing* in panel A, and *Data Sharing + Bing Default + Delayed Choice Screen* in Panel B. In all cases, data sharing reduces Google's market share by less than 0.5 percentage points. The effect is limited for the same reasons why equilibrium effects are small: economies of scale are small, and consumers show limited response to quality.

One caveat of the exercises in Table 10 is that our estimate of the demand response to the relevance of search results has a wide confidence interval. For that reason, Table A6 presents alternative equilibrium results where we take the largest demand response consistent with our estimates. Concretely, we use the lower bound of the 95 percent confidence interval for α on Table 7 rather than using our point estimate. We find larger equilibrium effects and effects of data sharing, but these effects are still small compared to those from counterfactuals that shut down demand-side frictions.

8 Conclusion

Google's large market share in web search is of ongoing concern to both antitrust authorities and regulators. This paper sheds light on this debate, focusing particularly on the role of browser defaults and economies of scale. Our results highlight that browser defaults are partially responsible for Google's large market share in web search. However, this effect does not arise only because of switching costs and users' inattention. Our findings show that consumers' lack of exposure to Bing—partly driven by browser defaults—is a key channel through which Google maintains a higher share than it would have absent any frictions. We also find that sharing Google's click-and-query data with Microsoft may only have a minor effect on market shares.

Our findings suggest that to significantly shift market shares, regulators must recognize search engines as experience goods and ensure that remedies expose consumers to alternatives. More broadly, our results provide a stark example of how overly pessimistic consumer beliefs about rivals can protect incumbent firms, rendering simple remedies ineffective.

References

Ackerberg, Daniel A, “Advertising, learning, and consumer choice in experience good markets: an empirical examination,” *International Economic Review*, 2003, *44* (3), 1007–1040.

Agte, Patrick, Claudia Allende, Adam Kapor, Christopher Neilson, and Fernando Ochoa, “Search and Biased Beliefs in Education Markets,” Technical Report, National Bureau of Economic Research 2024.

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow, “The welfare effects of social media,” *American economic review*, 2020, *110* (3), 629–676.

__ , Matthew Gentzkow, and Lena Song, “Digital addiction,” *American Economic Review*, 2022, *112* (7), 2424–2463.

Andersen, Steffen, John Y Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai, “Sources of inaction in household finance: Evidence from the danish mortgage market,” *American Economic Review*, 2020, *110* (10), 3184–3230.

Aridor, Guy, “Market definition in the attention economy: An experimental approach,” 2022.

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

Azevedo, Eduardo M, Alex Deng, José Luis Montiel Olea, Justin Rao, and E Glen Weyl, “A/B Testing with Fat Tails,” *Journal of Political Economy*, 2020, *128* (12), 4614–000.

Bajari, Patrick, Victor Chernozhukov, Ali Hortaçsu, and Junichi Suzuki, “The Impact of Big Data on Firm Performance: An Empirical Investigation,” in “AEA Papers and Proceedings,” Vol. 109 American Economic Association 2019, pp. 33–37.

Bellemare, Marc F., Jeffrey R. Bloem, and Noah Wexler, “The Paper of How: Estimating Treatment Effects Using the Front-Door Criterion*,” *Oxford Bulletin of Economics and Statistics*, 2024.

Blake, Thomas, Chris Nosko, and Steven Tadelis, “Consumer heterogeneity and paid search effectiveness: A large-scale field experiment,” *Econometrica*, 2015, *83* (1), 155–174.

Blake, Tom, Sarah Moshary, Kane Sweeney, and Steve Tadelis, “Price salience and product choice,” *Marketing Science*, 2021, *40* (4), 619–636.

Brot-Goldberg, Zarek, Timothy Layton, Boris Vabson, and Adelina Yanyue Wang, “The behavioral foundations of default effects: theory and evidence from Medicare Part D,” *American Economic Review*, 2023, *113* (10), 2718–2758.

Brynjolfsson, Erik, Avinash Collis, and Felix Eggers, “Using massive online choice experiments to measure changes in well-being,” *Proceedings of the National Academy of Sciences*, 2019, *116* (15), 7250–7255.

Bursztyn, Leonardo, Benjamin R Handel, Rafael Jimenez, and Christopher Roth, “When product markets become collective traps: The case of social media,” Technical Report, National Bureau of Economic Research 2023.

Carroll, Gabriel D, James J Choi, David Laibson, Brigitte C Madrian, and Andrew Metrick, “Optimal Defaults and Active Decisions,” *Quarterly Journal of Economics*, 2009, *124* (4), 1639–1674.

Chiou, Lesley and Catherine Tucker, “Search Engines and Data Retention: Implications for Privacy and Antitrust,” *Available at National Bureau of Economic Research 23815*, 2017.

Crawford, Gregory S and Matthew Shum, "Uncertainty and learning in pharmaceutical demand," *econometrica*, 2005, *73* (4), 1137–1173.

D.C. District Court, "United States of America et al. v. Google LLC," 2024, (Case No. 20-cv-3010 (APM)). Memorandum Opinion filed August 5, 2024.

Decarolis, Francesco, Muxin Li, and Filippo Paternollo, "Competition and Defaults in Online Search," *Available at SSRN 4660406*, 2023.

DellaVigna, Stefano and Ulrike Malmendier, "Paying Not to Go to the Gym," *American Economic Review*, 2006, *96* (3), 694–719.

Department of Justice, "United States vs Google LLC," 2020.

Dickstein, Michael J, "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice," 2021.

Dinielli, David, Fiona M Scott Morton, Katja Seim, Michael Sinkinson, Amelia Fletcher, Gregory S Crawford, Jacques Crémer, Paul Heidhues, Michael Luca, Tobias Salz et al., "Consumer Protection for Online Markets and Large Digital Platforms," 2023.

Duque, Omar Vásquez, "The Potential Anticompetitive Stickiness of Default Applications: Addressing Consumer Inertia with Randomization," *Available at SSRN 4077132*, 2022.

Ecommercedb, "Customer Journey in Online Shopping: Most Start on Search Engines," 2024.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz, "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," *American economic review*, 2007, *97* (1), 242–259.

Einav, Liran, Benjamin Klopack, and Neale Mahoney, "Selling Subscriptions," *Available at National Bureau of Economic Research 31547*, 2023.

Erdem, Tülin and Michael P Keane, "Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets," *Marketing science*, 1996, *15* (1), 1–20.

Ericson, Keith M Marzilli, "Consumer inertia and firm pricing in the Medicare Part D prescription drug insurance exchange," *American Economic Journal: Economic Policy*, 2014, *6* (1), 38–64.

European Commission, "CASE AT.40099, Google Android," 2018.

Farrell, Joseph and PAUL Klemperer, "Coordination and lock-in: Competition with switching costs and network effects," *Introduction to the series*, 2007, p. 1967.

Farronato, Chiara, Andrey Fradkin, and Chris Karr, "Webmunk: A New Tool for Studying Online Behavior and Digital Platforms," 2024.

Fowlie, Meredith, Catherine Wolfram, Patrick Baylis, C Anna Spurlock, Annika Todd-Blick, and Peter Cappers, "Default Effects and Follow-on Behaviour: Evidence from an Electricity Pricing Program," *The Review of Economic Studies*, 2021, *88* (6), 2886–2934.

Handel, Benjamin R, "Adverse selection and inertia in health insurance markets: When nudging hurts," *American Economic Review*, 2013, *103* (7), 2643–2682.

He, Di, Aadharsh Kannan, Tie-Yan Liu, R Preston McAfee, Tao Qin, and Justin M Rao, "Scale Effects in Web Search," in "Web and Internet Economics: 13th International Conference, WINE 2017, Bangalore, India, December 17–20, 2017, Proceedings 13" Springer 2017, pp. 294–310.

Heidhues, Paul, Alessandro Bonatti, L Elisa Celis, Gregory S Crawford, David Dinielli, Michael Luca, Tobias Salz, Monika Schnitzer, Fiona M Scott Morton, Michael Sinkinson et al., "More Competitive Search Through Regulation," *Yale J. on Reg.*, 2023, *40*, 915.

Ho, Kate, Joseph Hogan, and Fiona Scott Morton, "The impact of consumer inattention on insurer pricing in the Medicare Part D program," *The RAND Journal of Economics*, 2017, *48* (4), 877–905.

Hovenkamp, Erik, "The Competitive Effects of Search Engine Defaults," 2024. Available at SSRN: `https://ssrn.com/abstract=4647211` or `http://dx.doi.org/10.2139/ssrn.4647211`.

Imbens, Guido W, "Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics," *Journal of Economic Literature*, 2020, *58* (4), 1129–1179.

Israel, Mark, "Services as experience goods: An empirical examination of consumer learning in automobile insurance," *American Economic Review*, 2005, *95* (5), 1444–1463.

Klein, Tobias J, Madina Kurmangaliyeva, Jens Prüfer, Patricia Prüfer, and NN Park, "How Important Are User-Generated Data For Search Result Quality?," *Experimental evidence*, 2023.

Klemperer, Paul, "The competitiveness of markets with switching costs," *The RAND Journal of Economics*, 1987, pp. 138–150.

Lee, Kwok Hao and Leon Musolff, "Entry into two-sided markets shaped by platform-guided search," *Working Paper*, 2023.

Miller, Klaus M, Navdeep S Sahni, and Avner Strulov-Shlain, "Sophisticated Consumers with Inertia: Long-term Implications from a Large-scale Field Experiment," *Available at SSRN 4065098*, 2023.

Nelson, Phillip, "Information and consumer behavior," *Journal of political economy*, 1970, *78* (2), 311–329.

Ostrovsky, Michael, "Choice screen auctions," in "Proceedings of the 22nd ACM Conference on Economics and Computation" 2021, pp. 741–742.

Page, Larry, "2012 CEO Letter - Google Investor Relations," 2012. Accessed: 2024-11-10.

Patterson, Mark R, "Google and search engine market power," *Fordham Law Legal Studies Research Paper*, 2012, (2047047).

Pearl, Judea and Dana Mackenzie, *The book of why: the new science of cause and effect*, Basic books, 2018.

Schaefer, Maximilian and Geza Sapi, "Complementarities in Learning from Data: Insights from General Search," *Information Economics and Policy*, 2023, *65*, 101063.

Schmalensee, Richard, "Product differentiation advantages of pioneering brands," *The American Economic Review*, 1982, *72* (3), 349–365.

SEO Statistics, "124 SEO Statistics for 2024," 2024.

StatCounter, "Desktop Search Engine Market Share United States Of America," 2024.

__ , "Desktop vs. Mobile vs. Tablet Market Share: United States of America," 2024.

__ , "Search Engine Market Share Worldwide," 2024.

Stigler Committee on Digital Platforms, "Stigler Committee on Digital Platforms: Final Report," *Stigler Center News*, 2019.

The US House Judiciary Subcommittee on Antitrust, "Investigation of Competition in Digital Markets," 2020.

Tucker, Catherine, "Digital Data, Platforms and the Usual [Antitrust] Suspects: Network Effects, Switching Costs, Essential Facility," *Review of Industrial Organization*, 2019, *54* (4), 683–694.

UK Competition and Markets Authority, "Online Platforms and Digital Advertising Market Study," 2020.

Varian, Hal, "Is there a Data Barrier to Entry?," *Google*, 2015.

Varian, Hal R, "Position auctions," *international Journal of industrial Organization*, 2007, *25* (6), 1163–1178.

Online Appendix

Sources of Market Power in Web Search: Evidence from a Field Experiment

Hunt Allcott, Juan Camilo Castillo, Matthew Gentzkow, Leon Musolff, and Tobias Salz

Table of Contents

A Experimental Design Appendix

Figure A1: **Search Extension Window**

Notes: This figure shows the main window of Search Extension. Clicking the "Test Extension" button triggers an address bar search that the user cannot see and sends the anonymized data to our data repository.

B Experimental Results Appendix

Table A1: **Covariate Balance**

(a) **Baseline Google Users**

	(1) Control	(2) Active Choice	(3) Default Change	(4) Switch ($1)	(5) Switch ($25)	(6) Switch ($10 & CC)	(7) Switch ($10 & BC)	(8) Switch ($10 & CR)	(9) Switch ($10 & BR)	F-test p-value
Income ($000s)	52.30	62.48	59.64	59.07	61.94	53.80	52.18	55.96	57.59	0.40
College Degree	0.46	0.57	0.63	0.57	0.57	0.57	0.59	0.58	0.61	0.58
Male	0.33	0.44	0.52	0.47	0.46	0.42	0.39	0.47	0.48	0.13
Age	35.89	35.64	37.22	34.52	35.89	35.48	36.63	36.80	37.49	0.43
White	0.68	0.67	0.73	0.62	0.69	0.62	0.65	0.68	0.66	0.52

(b) **Baseline Bing Users**

	(1) Active Choice	(2) Switch ($10 & CC)	(3) F-test p-value
Income ($000s)	45.00	54.21	0.43
College Degree	0.47	0.45	0.82
Male	0.53	0.66	0.25
Age	35.13	37.61	0.37
White	0.66	0.63	0.81

Notes: Panels (a) and (b) present balance tests within the baseline Google user and baseline Bing user samples, respectively. Columns 1–9 (in Panel (a)) and 1–2 (in Panel (b)) present covariate means for each treatment group. The rightmost column presents the p-value of an F-test of a participant-level regression of that covariate on the treatment group indicators. The sample underlying this table includes all participants (inlcuding participants who did not stay with us until endline.)

Table A2: **Completion Rates**

(a) **Baseline Google Users**

	(1) Control	(2) Active Choice	(3) Default Change	(4) Switch Bonus ($1)	(5) Switch Bonus ($25)	(6) Switch Bonus ($10 & CC)	(7) Switch Bonus ($10 & BC)	(8) Switch Bonus ($10 & CR)	(9) Switch Bonus ($10 & BR)	(10) F-test p-value
Finished Survey 2	0.937	0.994	0.968	0.977	0.92	0.962	0.943	0.935	0.979	0.012
Kept Search Extension 2 weeks after Survey 2	0.873	0.964	0.943	0.93	0.857	0.92	0.874	0.906	0.936	0.009
Kept Search Extension 2 months after Survey 1	0.825	0.898	0.861	0.86	0.804	0.847	0.829	0.834	0.851	0.6

(b) **Baseline Bing Users**

	(1) Active Choice	(2) Switch Bonus ($10 & CC)	(3) F-test p-value
Finished Survey 2	0.921	0.921	1
Kept Search Extension 2 weeks after Survey 2	0.921	0.868	0.461
Kept Search Extension 2 months after Survey 1	0.842	0.868	0.748

Notes: Panels (a) and (b) present balanced attrition tests within the baseline Google user and baseline Bing user samples, respectively. Columns 1–9 (in Panel (a)) and 1–2 (in Panel (b)) present completion rates for each treatment group. The sample in each row is a strict subset of the row above. The first row presents the share of participants that completed Survey 2. The second row presents the share of participants that kept Search Extension installed for two days after completing Survey 2. The third row presents the share of participants that kept Search Extension installed for 14 days after completing Survey 2. The final row presents the share of participants that kept Search Extension installed eight weeks after completing Survey 1. The rightmost column presents the p-value of an F-test of a participant-level regression of completion indicators on the treatment group indicators.

Table A3: **Search Volume: Searches Per Day**

(a) **Baseline Google Users**

	(1) Control	(2) Active Choice	(3) Default Change	(4) Switch ($1)	(5) Switch ($25)	(6) Switch ($10 & CC)	(7) Switch ($10 & BC)	(8) Switch ($10 & CR)	(9) Switch ($10 & BR)	F-test p-value
t = 0	11.041	14.453	13.002	11.165	12.791	12.707	12.874	13.047	13.336	0.853
t = 1	10.036	14.220	12.564	12.995	15.157	16.194	14.463	14.177	14.905	0.141
t = 2	10.645	13.472	13.685	12.444	13.102	11.959	12.615	11.755	12.625	0.791
p-value (t = 0, t = 1)	0.402	0.665	0.708	0.308	0.089	0.000	0.037	0.085	0.018	-
p-value (t = 0, t = 2)	0.724	0.130	0.223	0.511	0.998	0.104	0.506	0.066	0.213	-

(b) **Baseline Bing Users**

	(1) Active Choice	(2) Switch ($10 & CC)	F-test p-value
t = 0	15.053	16.855	0.678
t = 1	14.278	14.872	0.872
t = 2	13.689	14.337	0.863
p-value (t = 0 vs t = 1)	0.576	0.258	-
p-value (t = 0 vs t = 2)	0.251	0.192	-

Notes: Panels (a) and (b) present the average number of searches per day on any search engine on each phase of the experiment for each of the treatment groups for baseline Google user and baseline Bing user samples, respectively. The phases are defined as follows: $t = 0$, $t = 2$, and $t = 3$ refer to the the days before Survey 1, the days between Survey 1 and Survey 2, and the days after Survey 2, respectively. The two bottom rows present the p-value of paired t-tests between the pre-Survey 1 average per day and the two other periods. The rightmost column presents the p-value of an F-test of a participant-level regression of the average search volume per day on the treatment group indicators.

Figure A2: **Initial Ratings of Google and Bing (Including Number of Ads)**

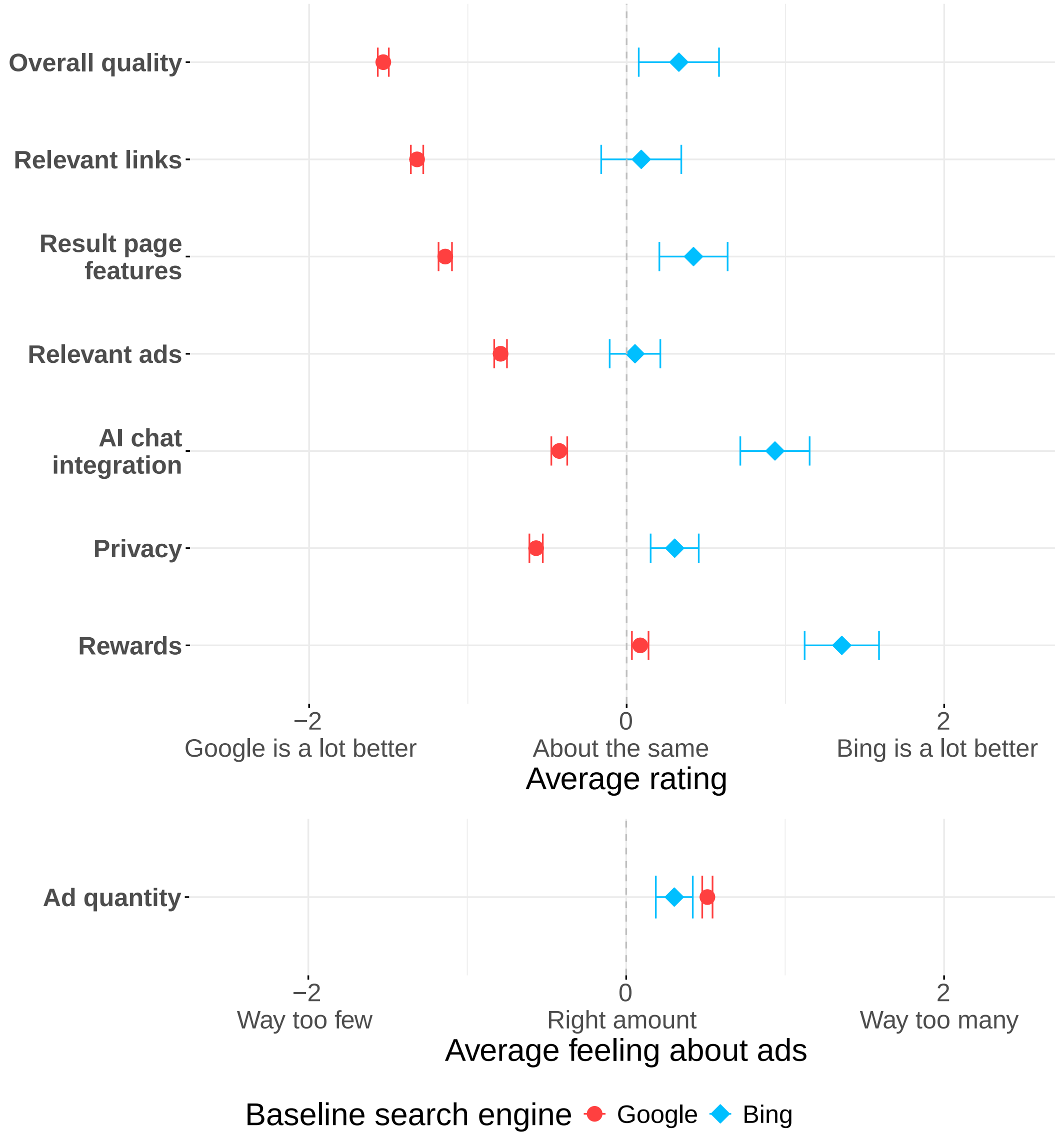

Notes: This figure presents average responses to the search engine rating questions for baseline Google and Bing users. The top rows present the average rating of Google and Bing on each reported dimension, in response to the following questions: "Overall, how would you rate the quality of Google relative to Bing?" and "How would you rate the quality of Google relative to Bing on the following dimensions?" Response options were "Bing is a lot better," "Bing is a little better," "They are about the same," "Google is a little better," and "Google is a lot better," coded as 2, 1, 0, -1, and -2, respectively. The bottom row presents the average response to the following question: "How do you feel about the number of ads on [baseline search engine used]?" Response options were "way too many," "too many," "right amount," "too few," and "way too few," coded as 2, 1, 0, -1, and -2, respectively. Whiskers indicate 95 percent confidence intervals.

Figure A3: **Why People Use Google or Bing**

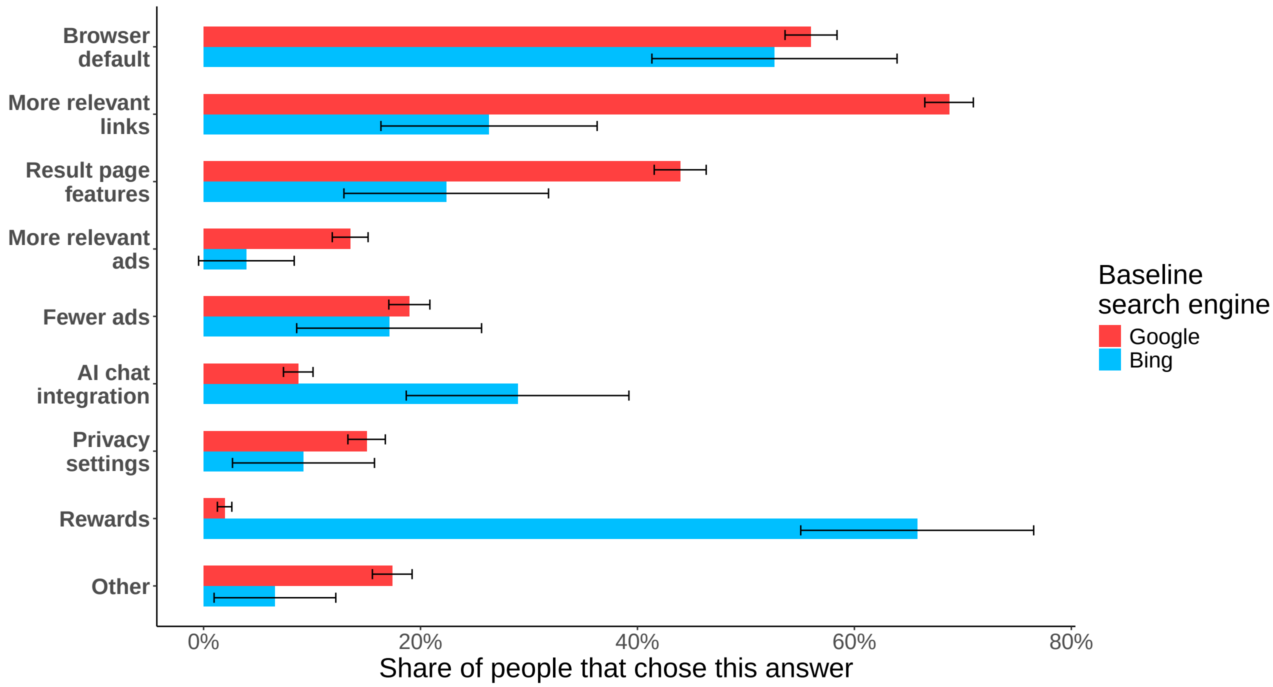

Notes: This figure presents the share of baseline Google and Bing users that chose each answer to the following question: "Why do you use [baseline search engine used] instead of [other search engine] for your searches on this web browser? Choose all that apply." Whiskers indicate 95 percent confidence intervals.

Figure A4: **Switch Treatment ($10) Search Engine Rating Change**

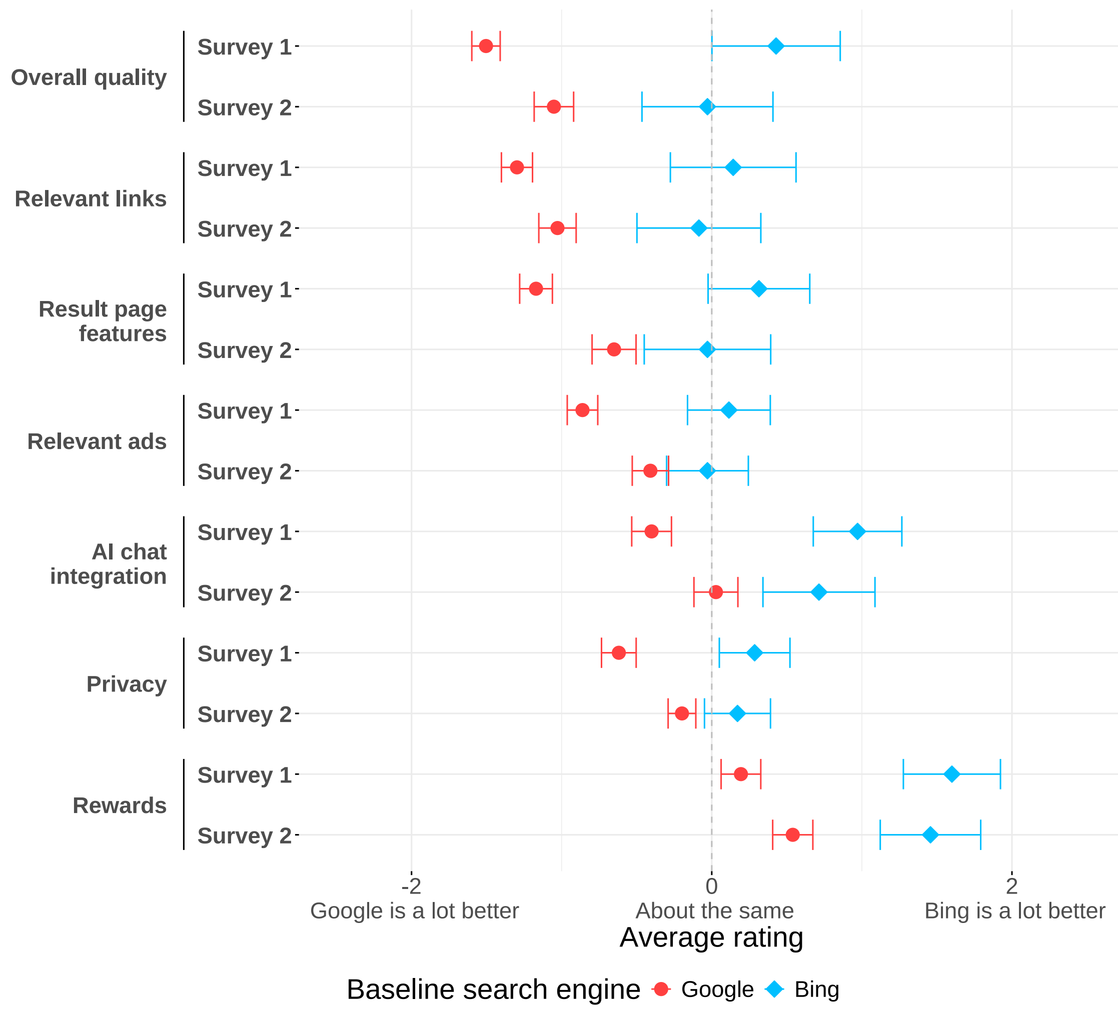

Notes: This figure presents average responses to the search engine rating questions for baseline Google and Bing users who were assigned to the $10 Switch Bonus Control group (S10CC). The responses to the rating questions were collected from both Survey 1 and Survey 2. The figure presents the average ratings of Google and Bing on each reported dimension, in response to the following questions: "Overall, how would you rate the quality of Google relative to Bing?" and "How would you rate the quality of Google relative to Bing on the following dimensions?" Response options were "Bing is a lot better," "Bing is a little better," "They are about the same," "Google is a little better," and "Google is a lot better," coded as 2, 1, 0, -1, and -2, respectively. Whiskers indicate 95 percent confidence intervals.

Figure A5: **Search Volume in Switch Bonus Groups**

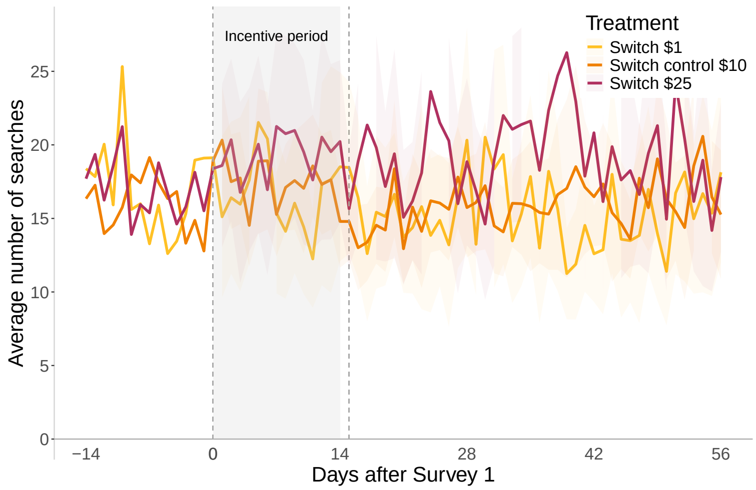

Notes: This figure presents the average number of searches on any search engine by participants of the Switch Bonus groups for each day of the experiment. The dashed vertical lines mark the dates of the two surveys. To arrive at daily market shares, we first compute the daily market share for each subject separately and then average shares across subjects.

Figure A6: **Baseline Ratings of Google vs. Bing**

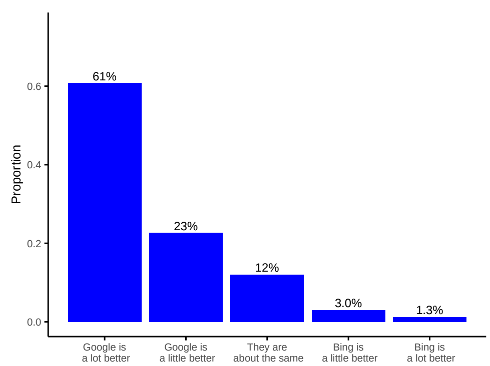

Notes: This figure presents the distribution of relative overall quality ratings reported on Survey 1. The survey question was, “Overall, how would you rate the quality of Google relative to Bing?”

C Demand Model Appendix

C.1 Identification Details

Distribution of idiosyncratic preferences and price response η. By comparing market shares of the Switch Bonus group at different payment offers during the incentivized period, we can identify the distribution of idiosyncratic preferences $S(\cdot)$ and the price response η, as shown in Figure 6. Consider a Switch Bonus user who is offered a price p to use search engine n after Survey 1. The utility from declining the Switch Bonus and staying with search engine d is

$$\zeta_d^* + \chi_{id} + \delta V_{i,t=2}(d). \tag{17}$$

The utility from accepting the Switch Bonus and switching the default to $j = -d$ is

$$\eta p + \tilde{\zeta}_{-d} + \chi_{i-d} + \delta V_{i,t=2}(-d). \tag{18}$$

Survey 1 tells participants that, regardless of their $t = 1$ choice, they will be guided through the choice screens on Survey 2 and consumers are therefore forced to pay the switching costs at time $t = 2$, even if they do not switch. Furthermore, we have assumed that consumers believe they know $\tilde{\zeta}_{-d}$ with certainty, so there is no perceived value from exploration. Therefore, $V_{i,t=2}(d) = V_{i,t=2}(-d)$ and the continuation values drop out from the comparison. Having this in mind, the consumer chooses $-d$ if

$$\eta p + \Delta\tilde{\zeta} + \Delta\chi_i > 0. \tag{19}$$

The modeled market share of search engine $-d$ at time $t = 1$ is thus

$$s_{-d,t=1}^{Sp} = S\left(\eta p + \Delta\tilde{\zeta}\right). \tag{20}$$

The parameters η and $\Delta\tilde{\zeta}$ simply play the role of a scale factor and a shifter, so we can rewrite this expression as

$$s_{-d,t=1}^{Sp} = H(p), \tag{21}$$

where $H(\cdot)$ is the cumulative density function of willingness to accept,$-1/\eta \cdot (\Delta\chi_i + \Delta\tilde{\zeta})$, a linear transformation of the idiosyncratic preferences term. Since the left hand side of Equation 21 is data, different price offers directly identify values of $H(\cdot)$ at different points. As is standard in discrete choice models, we normalize the mean and variance of $\Delta\chi_i$. This gives us the distribution $S(\cdot)$ of the normalized error $\Delta\chi_i$ from the shape of $H(p)$.[22]

[22]Technically, $H(\cdot)$ would be non-parametrically identified if we had a switch treatment for each price point p. In that case, $S(\cdot)$ would be non-parametrically identified up to a scale and location normalization. In practice, we have enough price points to determine that a log-normal fits our data well as illustrated in Figure 8: some participants are close to indifferent between Google

When $S(\cdot)$ is known, any two price points identify η, for instance:

$$\eta = \frac{S^{-1}(s^{S25}_{-d,t=1}) - S^{-1}(s^{S1}_{-d,t=1})}{25-1}$$

Perceived difference in quality $\Delta\tilde{\zeta}$. The perceived difference in quality is identified by market shares among Active Choice users, as depicted in Figure 6. Since these users made an active choice on Survey 1, we assume that their market shares from that point on are not driven by switching costs or inattention. Instead, they are entirely determined by perceived differences in quality.

Formally, for Active Choice users, the market share of the alternative search engine $-d$ is as in equation (8), with zero price difference Δp and without the switching cost $\sigma(1-\delta)$:

$$s^{A}_{-d,t\geq 1} = S(\Delta\tilde{\zeta}). \tag{22}$$

The relevant difference in quality is $\Delta\tilde{\zeta}$, capturing the fact that these users have not used search engine $-d$ and might thus have wrong perceptions about its quality. We can invert this equation to obtain the following expression for the perceived difference in quality $\Delta\tilde{\zeta}$:

$$\Delta\tilde{\zeta} = S^{-1}(s^{A}_{-d,t\geq 1}), \tag{23}$$

where $S^{-1}(x)$ is the inverse of the cumulative density function of $-\Delta\chi_i$.

Learning $\zeta^*_{-d} - \tilde{\zeta}_{-d}$. To identify learning, we compare the active choices made by Switch Bonus users after they had two weeks to learn about search engine $-d$ with the choices made by Active Choice users, who are not familiar with search engine $-d$ (see Figure 6).

When Switch Bonus subjects in the Search Extension Intervention Control group make their Survey 2 active choice, they have had time to learn ζ^*_{-d}, the true quality of search engine $-d$. They thus choose it if

$$\Delta\zeta^* + \Delta\chi_i > 0. \tag{24}$$

Assuming $\eta p_1 > \zeta^* - \tilde{\zeta}$, that is, that the Switch Bonus was large enough that all consumers who would choose Bing under perfect information were induced to try Bing, the market share of $-d$ is thus

$$s^{S10CC}_{-d,t\geq 2} = S\left(\Delta\zeta^*\right). \tag{25}$$

Comparing this expression with the market share for Active Choice users (equation 22) and rearranging gives the following expression:

$$\zeta^*_{-d} - \tilde{\zeta}_{-d} = S^{-1}\left(s^{S10CC}_{-d,t\geq 2}\right) - S^{-1}\left(s^{A}_{-d,t\geq 1}\right). \tag{26}$$

and Bing whereas others require large payments to abandon Google for two weeks.

Attention probability π. To explain how we identify π, we will first explain how the market share in the default group evolves.[23] Let $\tilde{s}^*$ be long-run share of the alternative search engine in the D group—i.e., the share after everybody who is not permanently inattentive has made an attentive choice. Also, let s^D_{-d*} be the share of the alternative search engine directly after treatment. Let $\tilde{\pi}$ be the rate at which participants who are not permanently inattentive become attentive in a given week.[24] With these definitions, the share of users who still use the alternative search engine in a given week is given by a weighted average of permanently inattentive users and users that become stochastically attentive; the users who become stochastically attentive converge away from s^D_{-d*} to $\tilde{s}^*$.

$$s^D_{-d,\text{week}=w} = \phi s^D_{-d*} + (1-\phi)\left[(1-\tilde{\pi})^w \cdot s^D_{-d*} + (1-(1-\tilde{\pi})^w)\cdot \tilde{s}^*\right]. \tag{27}$$

Intuitively, among those users who (i) are not permanently inattentive and (ii) would like to switch back, only a fraction $\tilde{\pi}$ actually switch back during a given week (which corresponds to half a period in our model). Therefore, the share of users switching away from the alternative search engine in a given week decays geometrically with a rate $\tilde{\pi}$, and we can identify that rate of decay from the following expression (which follows directly from equation (27)):

$$s^D_{-d,\text{week}=2} - s^D_{-d,\text{week}=1} = (1-\tilde{\pi})\left(s^D_{-d,\text{week}=1} - s^D_{-d,*}\right), \tag{28}$$

where s^D_{-d*} is the initial market share directly after treatment—that is, the fraction of users who switched to obtain our payment—and $s^D_{-d,\text{week}=w}$ represents the market share among the D group *at the end* of week w. Hence, we can derive the following expression for π directly as a function of market shares

$$\pi = 1 - \left(\frac{s^D_{-d,\text{week}=2} - s^D_{-d,\text{week}=1}}{s^D_{-d,\text{week}=1} - s^D_{-d,*}}\right)^2. \tag{29}$$

Switching costs σ and permanent inattention ϕ. As suggested by Figure 6, the switching cost σ and inattention parameter ϕ are jointly identified from the difference between the Active Choice and Control market share as well as the difference between the long run Default Change group market share and the Switch Bonus group market share. In particular, switching costs and inattention both create inertia: consumers are less likely to switch from the search engine they previously used, increasing the difference between the active choice and control group and the long-run default market share. However, as we will argue next, they are separately identified because σ affects each of these quantities symmetrically while inattention has a stronger effect on the long-run D group market share.

First, both types of inertia create a gap between the shares for Control users—who are subject to both forms of inertia—and for Active Choice users—who are not subject to either. The gap between those market

[23] Our explanation assumes that learning has not yet occurred, which means that the derived expressions are only correct for week one and two. This suffices for identification of the parameters.

[24] Given that π is defined for a two-week period, this is given by $\tilde{\pi} = 1-(1-\pi)^{1/2}$.

shares is

$$s^A_{-d,t\geq 0} - s^C_{-d,t\geq 0} = S\left(\Delta\tilde{\zeta}\right) - (1-\phi)S\left(\Delta\tilde{\zeta} - \sigma(1-\delta)\right), \tag{30}$$

which is increasing in both σ and ϕ.

Second, both types of inertia lead to higher market shares for the Default Change group after the incentive period. High switching costs and permanent inattention both imply that the geometric decay process that describes the market share over time will settle at a higher level. To state this formally, suppose for this section only that agents learn the true quality of the alternative search engine instantaneously after switching.[25] We now obtain an expression for $s^D_{-d,\infty}$, the value the market share $s^D_{-d,t}$ converges to as $t \to \infty$. Let $s^* = S(\Delta\zeta^* + \sigma(1-\delta))$ be the hypothetical long-run market share of the alternative search engine among D group users, assuming (i) everybody is attentive and (ii) users *have* learned its true quality. With these definitions, we can express the actual long-run market share of the alternative search engine as

$$s^D_{-d\infty} - s^{S10CC} = \phi s^D_{-d*} + (1-\phi)s^* = \phi s^D_{-d*} + (1-\phi)S(\Delta\zeta^* + \sigma(1-\delta)) - S(\Delta\zeta^*), \tag{31}$$

which is indeed increasing in ϕ and σ. To see why $s^D_{-d\infty}$ is the long-run market share note that everybody who is not permanently inattentive (fraction $1-\phi$) has made a choice and everybody else (fraction ϕ) is still stuck with the default.[26]

We now argue that although both moments (30 and 31) depend on switching costs and inattention, inattention has a much stronger effect on the latter. First, suppose there is no switching cost. In that case, $s^A_{-d,t\geq 0} - s^C_{-d,t\geq 0} = \phi S(\Delta\zeta^*) = \phi s^A_{-d,t\geq 0}$. Then note that permanent inattention affects both expressions as follows:

$$\frac{\partial}{\partial\phi}(s^A_{-d,t\geq 0} - s^C_{-d,t\geq 0}) = s^A_{-d,t\geq 0} \qquad \text{and} \qquad \frac{\partial}{\partial\phi}s^D_{-d\infty} = s^D_{-d*} - s^* \tag{32}$$

Thus, it affects the gap between the A and C groups to the extent that a lot of people in the active choice group want to use $-d$. As we saw in Section 4, few Chrome users want to use Bing , so permanent inattention will have little effect on the first expression. On the other hand, ϕ has a large impact on the long-run D share $s^D_{-d\infty}$ as long as (i) our treatment induces a large fraction of people s^D_{-d*} to use $-d$ in return for a payment, and (ii) many would not want to use it without payment, i.e. s^* is small. Conditions (i) and (ii) are both true in the data, as we saw in Section 4: over 75 percent of users switch in response to our payment, while the fraction who actually want to use is around 20 percent. Based on these numbers, we should expect the effect of permanent inattention on $s^D_{-d\infty}$ to be on the order of ten times larger than the effect on $s^A_{-d,t\geq 0} - s^C_{-d,t\geq 0}$.

Now suppose that there is no permanent inattention. Then $s^A_{-d,t\geq 0} - s^C_{-d,t\geq 0} = S(\Delta\zeta^*) - S(\Delta\zeta^* - \sigma(1-\delta))$

[25] We make this assumption only in this section to simplify the exposition. We otherwise maintain the assumption that people learn after fourteen days. The intuition extends to that case.

[26] We do not observe choices at $t = \infty$ as our sample ends after eight weeks, so our estimation (Section 5.2) uses the market share at the end of our experiment. Given our estimates from Section 5.3, the probability of paying attention after eight weeks is on the order of one thousandth, so the difference between these two expressions is negligible.

and $s^D_{-d\infty} = s^*$. Switching costs affect both expressions as

$$\frac{\partial}{\partial\sigma(1-\delta)}(s^A_{-d,t\geq 0} - s^C_{-d,t\geq 0}) = S'\left(\Delta\zeta^* - \sigma(1-\delta)\right) \qquad \text{and} \qquad \frac{\partial}{\partial\sigma(1-\delta)} s^D_{-d\infty} = S'\left(\Delta\zeta^* + \sigma(1-\delta)\right). \tag{33}$$

Therefore, the effect on both expressions should be roughly similar as long as the density of marginal users does not change too much.

The main takeaway from this analysis is that switching costs roughly have the same impact on $s^A_{-d,t\geq 0} - s^C_{-d,t\geq 0}$ and $s^D_{-d\infty}$, whereas permanent inattention has a much larger impact on $s^D_{-d\infty}$. This provides an argument why both parameters are separately identified.

Quality preferences α and ρ. Comparing the \$10 Switch Bonus at time $t = 2$ across the Ranking Degradation and Ad Blocking conditions identifies preferences for the components of quality. Similar to equation 25, the market share for a Switch Bonus user in the Search Extension Intervention $I \in \{RC, CA, RA\}$ is given by

$$s^{S10I}_{-d,t\geq 2} = S\left(\Delta\zeta^I\right). \tag{34}$$

where $\Delta\zeta^I$ is the quality implied by intervention I. Note that if $I = CC$—that is, if the user was assigned to the control group for both Ranking Degradation and Ad Blocking—then $\Delta\zeta^I = \Delta\zeta^*$.

Recall that search engine quality is given by $\zeta_j = \alpha a_j + \rho r_j + \xi_j$. The effect of Ranking Degradation on quality difference is $\zeta^{RC}_{-d} - \zeta^{CC}_{-d} = \zeta^{RA}_{-d} - \zeta^{CA}_{-d} = \rho\left(r^{RC}_{-d} - r^{CC}_{-d}\right)$. Thus, comparing Ranking Degradation relative to its control on $t \geq 2$ market shares (equations 25 and 34) and rearranging gives

$$\rho = \frac{S^{-1}\left(s^{S10RC}_{-d,t\geq 2}\right) - S^{-1}\left(s^{S10CC}_{-d,t\geq 2}\right)}{r^{RC}_{-d} - r^{CC}_{-d}}. \tag{35}$$

A similar expression can be obtained by comparing $s^{S10RA}_{-d,t\geq 2}$ and $s^{S10CA}_{-d,t\geq 2}$. Let r_j be defined in units of click through rates. Then, the right-hand side of equation (35) is observed in the data: it is the ratio of two treatment effects. Analogous equations also hold for the Ad Blocking condition, where we define a_j such that observed Bing ad load is $a_j = 1$.

C.2 Estimation Details

We now explain in detail the moments we use in our GMM procedure. The first set of moments are simply market shares: the baseline market share s_{-d0}, the Active Choice market share $s^A_{-d,t\geq 1}$, the market shares for the Switch Bonus group during the incentivized period at different prices $s^{S1}_{-d,t=1}$, $s^{S10CC}_{-d,t=1}$, and $s^{S25}_{-d,t=1}$, and the post-Survey 2 market shares of the Switch Bonus group under different interventions $s^{S10CC}_{-d,t\geq 2}$, $s^{S10RC}_{-d,t\geq 2}$, $s^{S10CA}_{-d,t\geq 2}$, and $s^{S10RA}_{-d,t\geq 2}$. To write out these nine moment conditions, we use m to index the moments that we target. For example, m can represent baseline choices for Chrome users, S10CC choices during the incentivized period for Edge users, or Active Choice choices at time $t \geq 2$ for Chrome users. We denote by $s_m(\theta)$ the market share predicted by our model for moment m when the model parameters are θ. We also define y_{mi} to be subject i's choice corresponding to moment m. Our first nine moment conditions take the form

$$g_{mi}(\theta) = y_{mi} - s_m(\theta), \qquad \mathbb{E}[g_{mi}(\theta^*)] = 0, \tag{36}$$

where θ^* is the vector of true parameters.

To identify the attention probability π, we exploit the market shares of the Default Change group right after Survey 1, after one week, and after two weeks. Rather than using these four market shares directly, we exploit our expression for the identification of π (equation 28). The moment condition that we use is

$$g_{mi}(\theta) = (1-\pi)^{1/2}(y_{m,i,\text{week}=1} - y_{m,i,*}) - (y_{m,i,\text{week}=2} - y_{m,i,\text{week}=1}), \qquad \mathbb{E}[g_{mi}(\theta^*)] = 0 \tag{37}$$

where $y_{m,i,*}$ denotes D group choices right after survey 1, and $y_{m,i,\text{week}=2}$ and $y_{m,i,\text{week}=2}$ denote D group choices at the end of weeks 1 and 2.

To identify switching costs and inattention, we need moments for $s^A_{-d,t\geq 0} - s^C_{-d,t\geq 0}$ and $s^D_{-d\infty}$. We already included moments corresponding to s_{-d0} (which is the same as $s^C_{-d,t\geq 0}$) and $s^A_{-d,t\geq 1}$, so we must include an additional moment for $s^D_{-d\infty}$. We use an empirical analogue of equation (31),

$$g_{mi}(\theta) = y_{m,i,\infty} - \phi y_{m,i,*} - (1-\phi)\left[\tilde{s}^*(\theta) + (1-\pi)^2(s^*(\theta) - \tilde{s}^*(\theta))\right], \qquad \mathbb{E}[g_{mi}(\theta^*)] = 0 \tag{38}$$

where $y_{m,i,*}$ denotes D group choices right after survey 1, and $y_{m,i,\infty}$ denotes D group choices after a long period has occured. In practice, we do not observe choices more than two months after the experiment starts, so our actual estimation uses an adjusted version of this moment that uses D group choices at the end of our experiment.[27] However, given our estimates from Section 5.3, the probability of not having paid attention after eight weeks is on the order of one thousandth, so the difference between these two moments is negligible.

There are two important issues we must deal with before computing these moment conditions and the GMM objective function. First, given the nature of our experiment, we don't observe all moments for every participant. For a participant that was randomized into S10CC, for instance, we observe the moments corresponding to S10CC choices at times $t = 1$ and $t \geq 2$, but we don't observe any of the Default Change or Active Choice choices. Second, we used different randomization probabilities for original Google and

[27] We now derive the expression for $s^D_{-d,\text{week}=8}$ that we use for estimation. After two weeks (that is, after learning) the fraction of people that would like to switch if attentive is no longer $s^D_{-d*} - \tilde{s}^*$ but $s^D_{-d*} - s^*$. The geometric decay process thus resembles equation 27, but it goes from s^D_{-d*} to s^* (and not from s^D_{-d*} to $\tilde{s}^*$). After accounting for the fraction $\left[1-(1-\tilde{\pi})^2\right](s^* - \tilde{s}^*)$ of users that switched back too early, we obtain the following expression for the market shares after week 2:

$$s^D_{-d,\text{week}=w>2} = \phi s^D_{-d*} + (1-\phi)\left[s^* + \left(1-(1-\tilde{\pi})^2\right)(s^* - \tilde{s}^*) + (1-\tilde{\pi})^w\left(s^D_{-d*} - s^*\right)\right].$$

Table A4: **Summary Statistics for Economy of Scale**

Description	Variable	Mean	Min	Median	Max
Number of Searches	n_{qt}	374	1	155	2,980
Dummy: First Result Clicked?	r_{qt}	0.23	0	0	1
Predicted Dummy: First Result Clicked?	$\hat{r}_{qt}$	0.31	0.02	0.18	1.14
Number of Observations			12,194,034		
Number of Queries			43,991		
Number of Top-Ranked URLs			244,136		

Bing users, so unconditional means would overweight Bing users in most of our treatments.

To address these issues, we think of our experiment using a potential outcomes setup. Hypothetically, for every moment m, there is a hypothetical realization of $g_{mi}(\theta^*)$. However, because of randomization, we do not observe many of these choices and thus, cannot compute the corresponding moments To address this issue, we rewrite our moment conditions in the form

$$\tilde{g}_{mi}(\theta) = w_{mi} \cdot g_{mi}(\theta), \qquad \mathbb{E}[\tilde{g}_{mi}(\theta^*)] = 0$$

where w_{mi} are weights that allow us to account for the fact that some of the moments $g_{mi}(\theta)$ are unobserved. Whenever $g_{mi}(\theta)$ is not observed, we simply set $w_{mi} = 0$ and $\tilde{g}_{mi}(\theta) = 0$. When $g_{mi}(\theta)$ is observed, we set w_{mi} to be the inverse of the (empirical) probability that we observe $g_{mi}(\theta)$ conditional on i's baseline search engine. Under these weights, it is still the case that $\mathbb{E}[\tilde{g}_{mi}(\theta^*)] = 0$ despite the fact that some of these choices are unobserved and that this occurs with different probabilities for baseline Bing and Google participants.

D Economies of Scale Appendix

D.1 Summary Statistics

Table A4 presents summary statistics for the click-and-query data we use to estimate economies of scale in data.

D.2 Descriptives

Before moving on to imposing a functional form, we use binscatters to present nonparametric plots of the relationship between the predicted click-through rate and the number of searches, controlling for query fixed effects. The left plot of Figure A7 exhibits the overall relationship, which seems to be roughly log-linear: each additional doubling in the number of searches leads to an about equal increase in click-through rate. The right plot separately analyzes this relationship for queries of differing popularity: while we find that the average level of click-through rate varies by query popularity, the relationship seems to be robustly

Figure A7: **Non-Parametric Estimates of Returns to Scale**

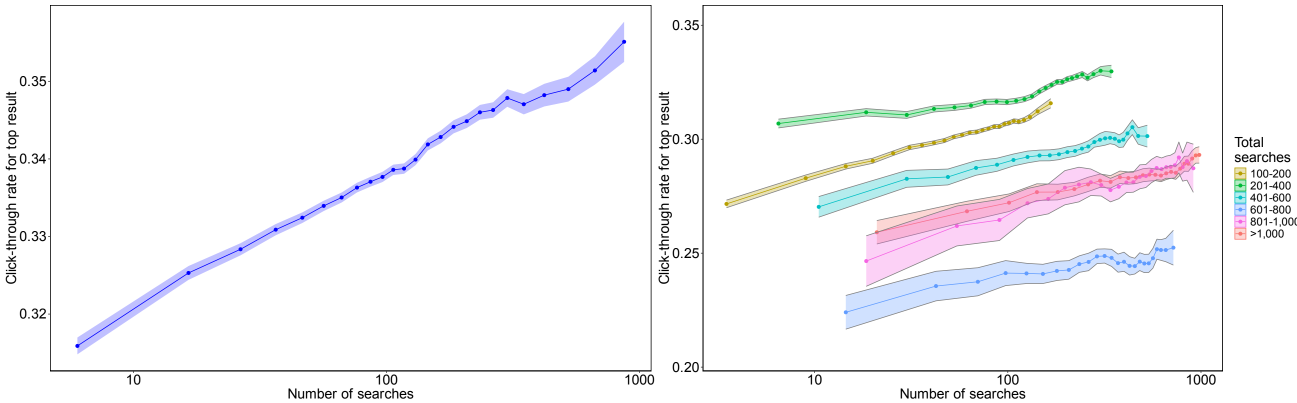

Notes: These graphs show the over-time relationship between the number of searches for a particular query and the rate of clicks on the top result (as disciplined by changes in the result shown first.) Formally speaking, we follow Section 6.1 and first regress the click-through rate on fixed effects for the top-ranked URL. We then use binsreg to analyze the relationship between the predicted click-through rate and the number of searches, controlling for query FE. The left plot exhibits the overall relationship, which seems to be roughly log-linear: each additional doubling in the number of views leads to an about equal increase in click-through rate. The right plot separately analyzes this relationship for queries of differing popularity: while we find that the average level of click-through rate varies by query popularity, the relationship seems to be robustly well-described as linear in the log of the number of views.

well-described as linear in the log of the number of views.

D.3 Implementation Details for Identification Argument

Our identifying assumption—that the causal effect of more data can only operate through the search ranking—allows us to apply the front door criterion (Pearl and Mackenzie, 2018; Imbens, 2020; Bellemare et al., 2024) to purge confounding variation in click-through rates, such as the observed secular trends. Focusing for now on a single query, the directed acyclical graph (DAG) in Figure A8 illustrates our strategy. In this graph, n_t refers to the number of searches for a given query by time t, ℓ_t to the the rankings of links on the results page at t, and r_t to the click-through rate (i.e., our measure of result relevance). Finally, u_t is an unobserved confounder that affects both the number of prior searches and the click-through rate, such as a changing composition of users over the lifetime of a query. For now, we assume these are all scalars (though they will not be in our eventual implementation.) The causal effect of additional prior searches (i.e., additional data) on the links served is given by γ and the causal effect of the links served on click-through rate is δ.

Regressing r_t directly on n_t would be biased by the confounding variation in u_t. The key insight from this graph is that by regressing r_t on ℓ_t after conditioning on n_t, one can isolate the causal variation in data that leads to changes in click-through rates. The intuition is that we stack many event studies like Figure 10: like in the event-study, we isolate the effect on CTR that comes with a change in the ranking. We then

Figure A8: **Directed Acyclic Graph Underlying Estimation Strategy**

$$\begin{tikzcd}
n_t \arrow[r, "\gamma"] & \ell_t \arrow[r, "\delta"] & r_t \\
& u_t \arrow[ul] \arrow[ur] &
\end{tikzcd}$$

Notes: This DAG represents how we identify the causal effect of increasing the number of searches (n_t) on the result relevance as measured by CTR (r_t) while accounting for an unobserved confounder (u_t) that could (e.g.) represent how different types of users arrive over the lifecycle of a query. The key to our identification strategy is the mediator ℓ_t, which represents the search engine's ranking of links. In particular, the search engine only learns from user data in the form of additional searches, and hence the relationship between n_t and ℓ_t is not confounded. Similarly, the relationship between ℓ_t and r_t is not confounded conditional on n_t.

regress the predicted value $\hat{r}_t = \hat{\delta} \cdot \ell_t$ on n_t and obtain the causal effect as the product $\gamma \cdot \delta$.

More formally, the identification challenge is that u_t may introduce a correlation between r_t and n_t, since $u_t \to n_t$ and $u_t \to r_t$. The regression

$$r_t = \alpha + \beta n_t + \varepsilon_t$$

would therefore lead to biased estimates of β, i.e. $\lim_{n \to \infty} \hat{\beta} \neq \gamma \times \delta$.

However, the confounder does not affect the ranking quality directly: there is no arrow from u_t to ℓ_t. Hence, the following regression (Bellemare et al., 2024, eq. 7)

$$\ell_t = \kappa + \gamma n_t + \omega_t$$

will yield an unbiased estimator $\hat{\gamma}$ of the effect of additional data (n_t) on ranking of links (ℓ_t). Similarly, we can run (Bellemare et al., 2024, eq. 8)

$$r_t = \lambda + \delta \ell_t + \phi n_t + \nu_t$$

to get an unbiased estimator $\hat{\delta}$ of the effect of the ranking of links (ℓ_t) on CTR (r_t). Multiplying together these two numbers we get an unbiased estimator $\hat{\gamma} \times \hat{\delta}$ of the effect of additional data on CTR.

In practice, we do not literally follow this recipe because of the complication introduced by the fact that the ranking of links ℓ_t is not a scalar. Keeping with it being a scalar, we now explain our alternative recipe. We predict CTR from ranking quality, i.e., we run

$$r_t = \lambda + \delta \ell_t + \phi n_t + \nu_t$$

and form a prediction $\hat{r}_t = \hat{\delta} \ell_t$ of CTR based just on the current ranking quality. Then we regress this prediction on the number of prior searches n_t, i.e. we effectively run

$$\hat{r}_t = \psi + \eta n_t + \varepsilon_t.$$

As $\hat{r}_t$ here is just ℓ_t multiplied by $\hat{\delta}$, this regression must yield $\hat{\eta} = \hat{\delta} \times \hat{\gamma}$, i.e., our estimator is numerically equivalent to what Bellemare and Bloem's recipe would find.

Now, we introduce our complication: ranking quality is measured by the identity of the top-ranked URL. Let $u(t)$ index the URL top-ranked as of search t (recall we still assume there is just one search term, so we do not need indices for search terms.) Then effectively ℓ_t is a vector of dummies: assuming there are U possible URLs that could be ranked first for this query,

$$\ell_t = (1(u(t) = 1), \ldots 1(u(t) = U))'.$$

This multidimensionality of ℓ_t makes the regression that Bellemare et al. (2024) propose hard to interpret and implement.

Still, our alternative way of first forming predictions of CTR works, and as argued above, in the scalar setting it would be exactly equivalent to employing the front-door criterion. Intuitively, we first project CTR on a fixed effect for the top-ranked URL while flexibly controlling for the number of searches a query has received so far. Subsequently, we use the fitted estimates from just the query-by-URL fixed effect in this regression as our dependent variable in estimating the relationship between searches and click-through rate. More formally, indexing queries by q and time by t, we first project CTR on a fixed effect for the top-ranked URL while controlling for a fixed effect for the number of searches a query has received so far, i.e.,

$$r_{qt} = \delta_{q,u(q,t)} + \eta_{n(q,t)} + \varepsilon_{qt}, \tag{39}$$

where $u(q,t)$ gives the index of the top-ranked result served on the search result page and $n(q,t)$ gives the number of searches that query q has seen by time t. As the regression includes a fixed effect $\eta_{n(q,t)}$, we are flexibly controlling for the number of searches a query has received so far. We then use the fitted estimates $\widehat{r}_{qt} = \delta_{q,u(q,t)}$ from just the query-by-URL fixed effect as our dependent variable in estimating the relationship between searches and click-through rate. Intuitively, these fitted estimates will capture systematic improvements in CTR that are driven by Bing changing the order in which it serves search results; by contrast, they will ignore changes due to pure temporal patterns (such as a secular trend in CTR.)

D.4 Taylor Expansion to Address HDFE in NLLS

Our main estimating equation (40) describes a non-linear relationship between the number of previous searches for a query and its click-through rate. However, the estimating equation also features a high-dimension fixed-effect, which is computationally challenging to estimate. To address this concern, this appendix develops a methodology that utilizes repeated Taylor expansions of an estimating equation to derive exact estimates of non-linear parameters in the presence of fixed-effects.

To begin with, we find initial estimates $(\hat{\beta}^0, \hat{\theta}^0)$ by regressing demeaned c_{qt} on demeaned $\frac{\beta}{1-\theta} n_{qt}^{1-\theta}$ (computing this term and then demeaning for any trial value of the parameters.) However, as the underlying regression is not linear, the resulting estimates from this exercise are possibly poor approximations to the true parameter values. To make progress, we turn the regression into a linear problem by utilizing a Taylor series expansion of 13 around initial estimates $(\hat{\beta}^0, \hat{\theta}^0)$. In particular, letting $\hat{\gamma}^0 = \frac{\hat{\beta}}{1-\hat{\theta}^0}$, we have

$$\hat{r}_{qt} = \alpha_q + \gamma n_{qt}^{1-\hat{\theta}^0} + \gamma(\hat{\theta}^0 - \theta)\log(n_{qt}) n_{qt}^{1-\hat{\theta}^0} + O(\theta - \hat{\theta}^0)^2 + O(\gamma - \hat{\gamma}^0)^2 + \varepsilon_{qt}.$$

As this equation is linear in easily constructed regressors $n_{qt}^{1-\hat{\theta}^0}$ and $\log(n_{qt}) n_{qt}^{1-\hat{\theta}^0}$, it can be estimated while correctly accounting for the FE α_q, thus yielding new estimates $(\hat{\beta}^1, \hat{\theta}^1)$. We can then form a new Taylor expansion around those estimates, yielding$(\hat{\beta}^2, \hat{\theta}^2)$ and so on. We iterate until convergence, and obtain standard errors via block-boostrapping (resampling at the query-level.)

D.5 Effect of Data on Result Relevance

We now use our estimates in Table 8 to anticipate by how much Bing's CTR would increase if it were to obtain additional data, which could come either from an increase in its market share or from regulatory provisions that require the sharing of click and query data. Note that all estimates in this subsection take a partial equilibrium approach, i.e., they do not consider the effect that an improvement in Bing's CTR may have on its market share and the feedback loop that could potentially result from this effect. When moving to the full model below, we will take into account this flywheel.

Suppose Bing was to obtain an additional 1,000 searches on each query. This would result in an increase of CTR from 23.5 percent to 25.0 percent, an increase of 1.55 percentage points. We can see in Figure A9 that this increase mostly comes from an improvement in serving results on tail queries. Similarly, what if Bing multiplied it's market share by 4.28, making it roughly equal to Google's market share? In this case, our estimates imply that Bing's CTR would increase from 23.5 percent to 24.8 percent, an increase of 1.29 percentage points.

More generally, we can use our parameter estimates to calculate counterfactual average click-through rates on Bing if its market share were to be multiplied by λ: the new click-through rate c' after such an increase in market share is given as a function of the old click-through rate by

$$r' = \alpha - \frac{\beta}{1-\theta} + \lambda^{1-\theta} \cdot \left(r - \alpha + \frac{\beta}{1-\theta}\right)$$

D.6 Cross-Query Learning

A complication that we do not address in the main text is that learning how to rank results on query q may not be limited to using data from query q – customer behavior on impressions on related queries q' are also helpful. Denote the set of all queries (including the focal query q) as Q, and define a distance metric $d(q, q')$

Figure A9: **Marginal Effect of Additional Impressions on Bing CTR by Query Popularity**

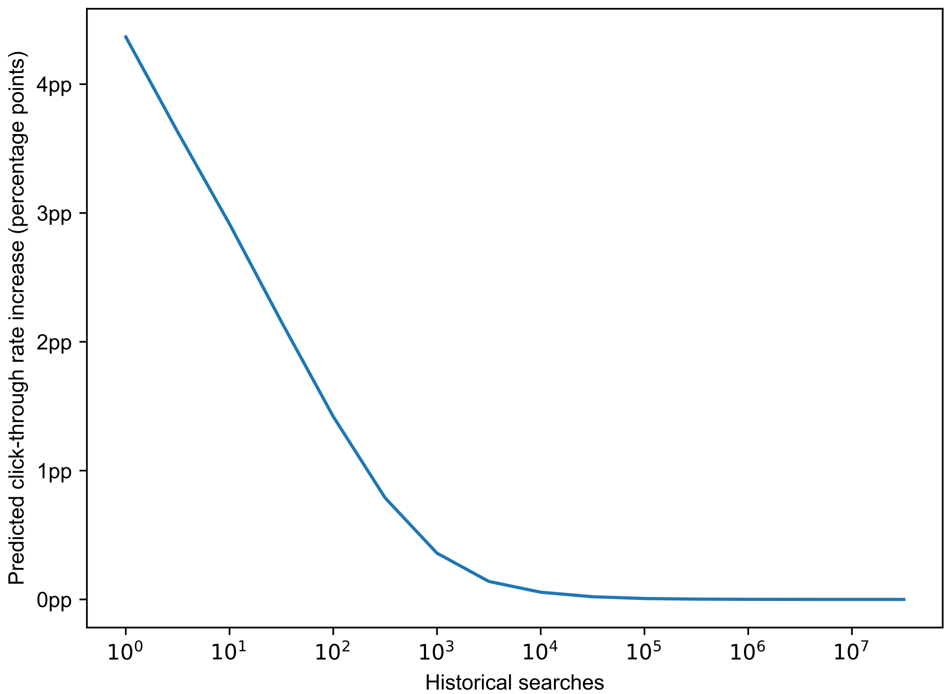

Notes: This figure shows the effect (in percentage points) on CTR of increasing the number of searches that Bing observes for each query by 1,000. The resulting improvement in CTR is concentrated on tail queries, which benefit the most from additional data; queries that already had large quantities of data improve less.

that measures the distance between any focal query q and potentially related query q'. The distance between any query and itself is zero, i.e., $d(q,q) = 0$ for all q. Similar to in the main text, we will assume that the CTR on query q is given

$$\hat{r}_{qt} = \alpha_q + \beta \frac{1}{1-\theta} \left(\sum_{q' \in Q} f(d(q,q');\gamma) \times n_{q't} \right)^{1-\theta} + \varepsilon_{qt} \tag{40}$$

As before, α_q is a fixed effect that captures the fact that different queries may have different baseline CTR. More importantly, β is our measure of the value of gathering additional data, and θ measures economies of scale, i.e., the speed at which this value declines with the amount of data already gathered. In particular, $\theta \approx 0$ implies linear returns to additional data, $\theta \approx 1$ implies logarithmic returns, and $\theta >1$ implies worse-than-logarithmic returns.

Finally, $f(\cdot)$ is a function (parameterized by γ) that maps the distance $d(q,q')$ between a focal query q and a related query q' into a monotonically declining weight. The speed at which these weights decay as we consider more and more distant queries play an important role in the economies of scale: if the weight decays only slowly, it may not matter if a search engine has never seen a query before as it can apply it's learnings from other, related queries. If the weight decays quickly, on the other hand, then not having seen a particular query before would be a strong disadvantage in trying to serve its results. Given our limited data, we will parameterize

$$f(d) = \exp(-\exp(\gamma)d), \tag{41}$$

so that $\gamma = -\infty$ corresponds to no decay with distance (i.e., all views on all queries matter to CTR on any focal query) and $\gamma \to \infty$ corresponds to the case of no cross-query learning (but still allows views on the focal query to matter as $\lim_{\gamma\to\infty} \exp(-\exp(\gamma) \times 0) = 1$.)

To estimate (40), we supplement the data on new queries by obtaining, for each query in the original dataset, similar data on the 50 other queries most related to the original query, as reported by Bing's internal metrics. We emphasize that this means we have data only on the most related queries in Q; to the extent that there is little cross-query learning, we would expect this to not bias our results as searches for less related queries would not yield additional learning on Bing's side. The supplemental dataset contains, aggregated to the daily level, the total number of searches for and clicks on results pages of each of these related queries, again between 2022-01-24 and 2023-01-23. We note that, by construction, these related queries are not necessarily new; hence, we also obtain the total number of impressions between 2021-01-24 and 2022-01-23 for the related queries (this number is by definition zero for the focal queries.) Whenever we consider a running tally of searches in our estimation below, we consider the period between 2022-01-24 and 2023-01-23 (when we see all searches) and add to the searches that have ocurred by any given date during this period the searches that occurred in the prior year, i.e., from 2021-01-24 and 2022-01-23. However, we cannot account for searches even further in the past due to Bing's retention policy for query data. Finally, we have access to Bing's internal distance measure between the related queries and the original focal query.

As before, we use the fitted estimates $\widehat{r}_{qt}$ using just the $\delta_{q,u(q,t)}$ fixed effect from regression (9) as our dependent variable in the estimation of (40). Computationally speaking, we obtain our estimates of (40) via a non-linear least squares procedure and standard errors from a block bootstrap (where a block is a focal query.) As our non-linearity correction from Section 6.2 did not yield substantively different estimates there, we avoid implementing a more complicated version of this procedure in this robustness check and simply report the parameters estimated via our demeaned non-linear least squares procedure, noting that these estimates should be interpreted with caution.

As before, we calibrate the intercept α such that the our average predicted CTR matches that from our experiment. However, a complication emerges: to take this average, we need to know for each query in the query frequency distribution how many views there are on related queries. However, we also require information on the number of views on related queries. While we have this information for the new queries (used in estimation above), we do not have this information for all queries that Bing sees. Hence, we need to predict the value of the term in parentheses, i.e., $\sum_{q' \in Q} f(d(q,q');\gamma) \times searches_{q't}$ from just the number of views on the focal query. We use the model

$$log(\sum_{q' \in Q} f(d(q,q');\gamma) \times searches_{q't}) = \beta_0 + \beta_1 log(searches_q) + u_q \tag{42}$$

We fit this equation on our sample of new queries (for which we observe views on related queries), and find $\hat{\beta}_0$ =0.0509(0.0058) and $\hat{\beta}_1$ =0.9951(0.0010) with R^2 =0.94, suggesting that we can predict this quantity very well. We can thus use

$$r_{qt} = \alpha_q + \beta \frac{1}{1-\theta} \left(\exp(\beta_0 + \beta_1 log(n_q))\right)^{1-\theta} + \varepsilon_{qt}$$

to predict CTR from just an observation of the number of views on a particular (focal) query. This allows us to calibrate α.

We exhibit our results in Table A5. Most importantly, we still find returns that are essentially logarithmic, though once taking into account spillovers, the returns are slightly more convex than logarithmic (i.e., $\hat{\theta} < 1$). Furthermore, we can strongly reject the null hypothesis that additional searches have no impact on performance (i.e., β is significantly different from zero.) Finally, our estimate of γ suggests a limited role of cross-query learning. This is illustrated by Figure A10, which plots the implied weight of searches on related queries against their distance from the focal query. In particular, the horizontal axis measures the distance to the focal query in units of Bing's internal distance metric; these units are restricted to lie between zero (only assigned for identical queries) and two (practically never assigned.) The solid black line indicates the weight that our estimates imply for views on a query at a certain distance from a focal query: for instance, at distances of 0.01 our estimates imply a weight of about 0.4, suggesting that each search for a related query at this distance is worth about 40 percent of a search for the original query when it comes to learning how to rank search results. As query distances are hard to interpret, we exhibit the distribution

Table A5: **Economy of Scale Estimates**

Description	Parameter	Estimate	SE
Click-through rate at inception	α	0.1744	-
Value of additional data	β	0.0056	(0.0007)
Shape of returns from data	θ	0.9272	(0.0292)
Relative weight on related queries	γ	3.8327	(0.6659)

Notes: This table provides the estimates of the parameters in (40), obtained via non-linear least squares. Standard errors are from a block-bootstrap clustered at the focal query level.

of distances between focal queries and their top-related query (in blue) or their tenth-most related query (in red). We can see that it is rare for queries to have a related query at distance low enough to be assigned a significant weight.

According to the estimates in Table A5, if Bing were to increase its market share by multiplying it by 4.28x, its CTR would improve from 23.50 percent to 24.99 percent, an improvement of 1.49 percentage points. As expected, this increase is slightly larger than the 1.29 percentage points that we found for an increase in market share by multiplying it by 4.28x in the main text. In other words, accounting for cross-query spillovers slightly raises our estimates of the importance of economies of scale, but does not lead to any changes in qualitative conclusions.

E Counterfactuals Appendix

Consider the utility of agent i in some counterfactual $\mathscr{C}$. The difference in the user's perceived utilities—the utility that drives choices—can be written as

$$\Delta u_{i,\mathscr{C}} = \Delta v_{\mathscr{C}} + \Delta b_{\mathscr{C}} + \Delta \chi_i,$$

where $\Delta v_{\mathscr{C}}$ denotes differences in true mean utilities, $\Delta b_{\mathscr{C}}$ denotes additional differences due to misperceptions, and $\Delta \chi_i$ denotes differences in idiosyncratic preferences. To illustrate these terms, we now consider what they look like in the Status Quo. The term for true utilities is $\Delta v_{\mathscr{C}} = \Delta\zeta - \sigma(1-\delta)$ to account for the difference in the quality of the search engine and for the switching cost. The bias term is $\Delta b_{\mathscr{C}} = \zeta^*_{-d} - \zeta_{-d}$ since the user is not aware of the true quality of the alternative search engine.

The market share of search engine n is given by

$$s_{-d,\mathscr{C}} = S(\Delta v_{\mathscr{C}} + \Delta b_{\mathscr{C}}),$$

where $S(\cdot)$ is the CDF of the difference in idiosyncratic preferences.

We now derive an expression for consumer surplus. In the absence of any misperceptions, the consumer surplus relative to the utility of the original search engine is given by $\frac{1}{\eta}V(\Delta v_{\mathscr{C}})$, where $V(x) = \int_{-\infty}^{x} S(x')dx'$.

Figure A10: **Cross-Query Learning**

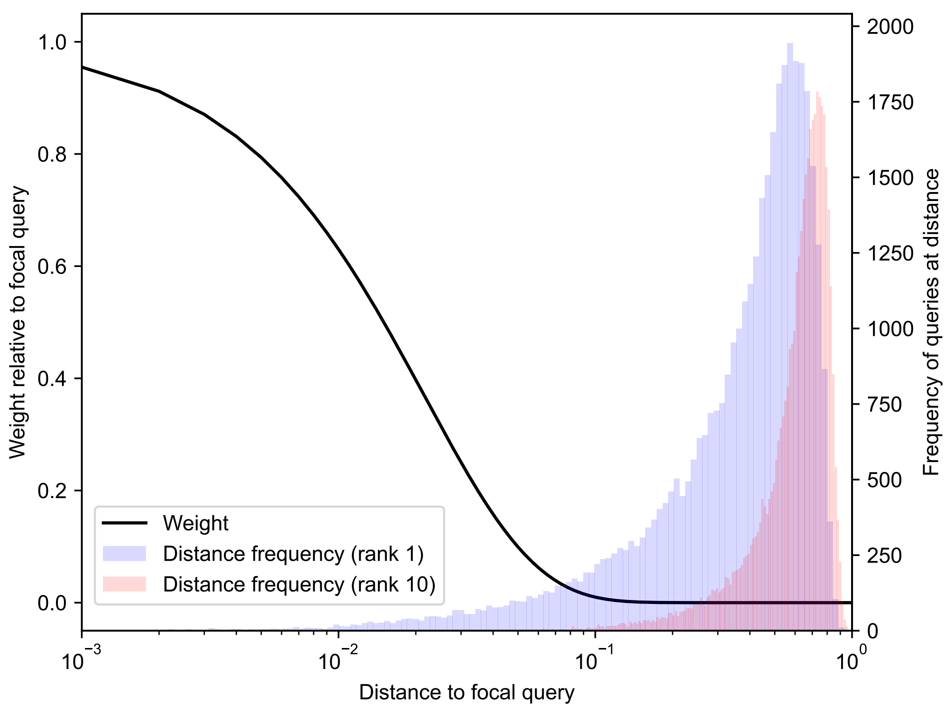

Notes: This graph illustrates the (limited) extent of cross-query learning implied by our parameter estimates. The horizontal axis measures the distance to the focal query in units of Bing's internal distance metric; these units are restricted to lie between zero (only assigned for identical queries) and two (practically never assigned.) The solid black line indicates the weight that our estimates imply for views on a query at a certain distance from a focal query: for instance, at distances of 0.01 our estimates imply a weight of about 0.4, suggesting that each view on a related query at this distance is worth about 40 percent of a view on the original query when it comes to learning how to rank search results. As query distances are hard to interpret, we exhibit the distribution of distances between focal queries and their top-related query (in blue) or their tenth-most related query (in red). We can see that it is rare for queries to have a related query at distance low enough to be assigned a significant weight.

The consumer surplus relative to the utility of the original search engine *in the Status Quo* is given by

$$CS_{\mathscr{C}} = \frac{1}{\eta}\left[V(\Delta v_{\mathscr{C}} + \Delta b_{\mathscr{C}}) - \Delta m_{\mathscr{C}} \cdot S(\Delta v_{\mathscr{C}} + \Delta b_{\mathscr{C}}) + (v_{o,\mathscr{C}} - v_{o,SQ})\right]$$

The first term in this expression gives the consumer surplus, relative to the utility of the original search engine in $\mathscr{C}$, if users's true utility is indeed described by their perceived utility. The second term is an adjustment due to biases that lead to suboptimal choices. For the share $S(\Delta v_{\mathscr{C}} + \Delta b_{\mathscr{C}})$ of users that choose the alternative search engine, their true utility is lower (or higher) to the extent that they have biases, $-\Delta b_{\mathscr{C}}$. The final term accounts for the fact that we want to measure utility relative to the Status Quo to be able to compare consumer surplus across counterfactuals. We therefore must adjust consumer surplus by the degree to which utility at the anchor point–the original search engine–changed relative to the status quo, $v_{d,\mathscr{C}} - v_{d,SQ}$.

E.1 Direct Effects

We now give expressions for $\Delta v_{\mathscr{C}}$, $\Delta b_{\mathscr{C}}$, and $(v_{d,\mathscr{C}} - v_{d,SQ})$ in each of our counterfactuals. Note that the Data Sharing and Data Sharing + Choice Screen counterfactuals are only relevant in equilibrium, since they involve a change in the quality of search engines that arises from the use of data.

Status Quo For the status quo, $\Delta v_{SQ} = \Delta\zeta - \sigma(1-\delta)$ to account for differences in utilities and for welfare-relevant switching costs. The bias term $\Delta b_{SQ} = \tilde{\zeta}_{-d} - \zeta^*_{-d}$ accounts for misperceptions and for the fraction of switching costs that is not welfare-relevant. Trivially, $(v_{d,SQ} - v_{d,SQ}) = 0$.

No Frictions Since there are no switching costs, $\Delta v_{NF} = \Delta\zeta$. Since there are no biases, $\Delta b_{NF} = 0$. And the true quality of the original search engine is unchanged, so $(v_{d,NF} - v_{d,SQ}) = 0$.

Active Choice Since there are no switching costs, $\Delta v_{CS} = \Delta\zeta$. Customers still have misperceptions about search engines, so $\Delta b_{NF} = \tilde{\zeta}_{-d} - \zeta^*_{-d}$. And the true quality of the original search engine is unchanged, so $(v_{d,NF} - v_{d,SQ}) = 0$.

Correct Perceptions True utilities do not change, so $\Delta v_{CP} = \Delta\zeta - \sigma(1-\delta)$. There are no misperceptions, so $\Delta b_{CP} = 0$. And the true quality of the original search engine is unchanged, so $(v_{d,CP} - v_{d,SQ}) = 0$.

Choice Screen Edge users behave as in the Status Quo. Chrome users behave as in the Active Choice counterfactual.

Bing Default Edge users behave just as in the Status Quo. Among Chrome users, switching costs are still present but they go the other way around, so $\Delta v_{BD} = \Delta\zeta + \sigma(1-\delta)$. We assume there are no misperceptions about search engines, so $\Delta b_{BD} = 0$. Finally, the true quality of the original search engine changes because it is now subject to switching costs, so $(v_{d,BD} - v_{d,SQ}) = -\sigma(1-\delta)$.

Bing Default + Delayed Choice Screen During the first two weeks, the market behaves just as in the Bing Default counterfactual. Starting on week 3, Edge users behave as in the Status Quo. Chrome users behave as in the Correct Perceptions counterfactual. We compute market shares and consumer surplus over a span of four years: we give weight 2/208 to the first two weeks, and we give weight 206/208 to the market after week 3.

Bing Payments Utilities change due to payments, so $\Delta v_{BP} = \eta \Delta p + \Delta\zeta - \sigma(1-\delta)$. Customers still have the same biases as in the status quo, so $\Delta b_{CP} = \tilde{\zeta}_{-d} - \zeta^*_{-d}$. The true quality of the original search engine is unchanged, so $v_{d,NF} - v_{d,SQ} = 0$ for Chrome users. For Edge users, we must account for the fact that the utility of using Bing increased by \$10, so $v_{d,BP} - v_{d,SQ} = \eta\Delta p$.

E.2 Equilibrium Effects

To account for equilibrium effects, we must account for the fact that true qualities ζ_j are now a function of the share of people using search engines. Suppose that a share $\bar{s}_j$ of people use search engine *j across all browsers*. Following our economies of scale model, they are given by

$$\zeta_j(\bar{s}_j) = \rho\left[\alpha - \frac{\beta}{1-\theta} + \left(\frac{\bar{s}_j}{\bar{s}_{j,SQ}}\right)^{1-\theta}(\hat{r}_j - \alpha + \frac{\beta}{1-\theta})\right] + \xi_j$$

We can thus derive the following expressions for $\Delta v_{\mathscr{C}}$, $\Delta b_{\mathscr{C}}$, and $(v_{o,\mathscr{C}} - v_{o,SQ})$ in equilibrium.

No Frictions Since there are no switching costs, $\Delta v_{NF} = \Delta\zeta = \zeta^*_{-d}(\bar{s}_{-d,NF}) - \zeta^*_d(1-\bar{s}_{-d,NF})$. Since there are no biases, $\Delta b_{NF} = 0$. The true quality of the original search engine changes with the new market shares, so $v_{d,NF} - v_{d,SQ} = \zeta^*_d(1-\bar{s}_{-d,NF}) - \zeta^*_d(1-\bar{s}_{-d,SQ})$.

Active Choice Since there are no switching costs, $\Delta v_{AC} = \Delta\zeta = \zeta^*_{-d}(\bar{s}_{-d,AC}) - \zeta^*_d(1-\bar{s}_{-d,AC})$. Customers still have misperceptions about search engines, so $\Delta b_{AC} = \tilde{\zeta}_{-d} - \zeta_{-d}(\bar{s}_{-d,AC})$. The true quality of the original search engine changes with the new market shares, so $v_{d,AC} - v_{d,SQ} = \zeta^*_d(1-\bar{s}_{-d,AC}) - \zeta^*_d(1-\bar{s}_{-d,SQ})$.

Correct Perceptions True utilities change due to new market shares, so $\Delta v_{CP} = \zeta^*_{-d}(\bar{s}_{-d,CP}) - \zeta^*_d(1-\bar{s}_{-d,CP}) - \sigma(1-\delta)$. There are no misperceptions, so $\Delta b_{CP} = 0$. The true quality of the original search engine changes with the new market shares, so $v_{d,CP} - v_{d,SQ} = \zeta^*_d(1-\bar{s}_{-d,CP}) - \zeta^*_d(1-\bar{s}_{-d,SQ})$.

Correct Perceptions + Data Sharing All expressions are the same as in Correct Perceptions, except that the quality of Bing is given by $\zeta^*_{-d}(1)$ instead of $\zeta^*_{-d}(\bar{s}_{-d,CP})$.

No Frictions + Data Sharing All expressions are the same as in No Frictions except that the quality of Bing is given by $\zeta^*_{-d}(1)$ instead of $\zeta^*_{-d}(\bar{s}_{-d,NF})$.

Choice Screen Edge users behave as in the Status Quo, and Chrome users behave as in the Active Choice counterfactual. Qualities must be adjusted to account for equilibrium effects: they are now $\zeta^*_{-d}(\bar{s}_{-d,CS})$ and $\zeta^*_d(1-\bar{s}_{-d,CS})$.

Bing Default Edge users behave just as in the status quo, but qualities do change because of economies of scale: $\Delta v_{BD,E} = \Delta\zeta - \sigma(1-\delta)$, $\Delta b_{BD,E} = \tilde{\zeta}_{-d} - \zeta^*_{-d}(\bar{s}_{-d,BD})$, and $(v_{d,SQ} - v_{d,SQ}) = 0$. The true quality of Bing changes: $(v_{d,BD,E} - v_{d,BD,E}) = \zeta^*_d(1-\bar{s}_{-d,BD}) - \zeta^*_d(1-\bar{s}_{-d,SQ})$. For Chrome users, switching costs are still present but they go the other way around::$\Delta v_{BD,C} = \zeta^*_{-d}(\bar{s}_{-d,BD}) - \zeta^*_d(1-\bar{s}_{-d,BD}) + \sigma(1-\delta)$. There are no misperceptions so $\Delta b_{BD,C} = 0$, and the true quality of the original search engine changes due to the change in market shares and and because it is now subject to switching costs so $(v_{d,BD,C} - v_{d,BD,C}) = \zeta^*_d(1-\bar{s}_{-d,BD}) - \zeta^*_d(1-\bar{s}_{-d,SQ}) - \sigma(1-\delta)$.

Bing Default + Delayed Choice Screen During the first two weeks, the market behaves just as in the Bing Default counterfactual. Starting on week 3, Edge users behave as in the Status Quo. Chrome users behave as in the Perfect Information counterfactual: $\Delta v_{DCS,C} = \zeta^*_{-d}(\bar{s}_{-d,DSC}) - \zeta^*_d(1-\bar{s}_{-d,DSC}) - \sigma(1-\delta)$, $\Delta b_{DSC,C} = 0$, and $(v_{d,DSC,C} - v_{d,SQ}) = \zeta^*_d(1-\bar{s}_{-d,DSC}) - \zeta^*_d(1-\bar{s}_{-d,SQ})$. We compute market shares and consumer surplus using the same weights as in the direct effects counterfactual.

Bing Payments Utilities change due to payments. Thus, $\Delta v_{BP} = \eta\Delta p + \zeta^*_{-d}(\bar{s}_{-d,BP}) - \zeta^*_d(1-\bar{s}_{-d,BP}) - \sigma(1-\delta)$ for Chrome users and $\Delta v_{BP} = -\eta\Delta p + \zeta^*_{-d}(\bar{s}_{-d,BP}) - \zeta^*_d(1-\bar{s}_{-d,BP}) - \sigma(1-\delta)$ for Edge users. Customers still have the same biases as in the status quo, so $\Delta b_{BP} = \tilde{\zeta}_{-d} - \zeta^*_{-d}(\bar{s}_{-d,BP})$. The true quality of the original search engine changes with the new market shares, so $v_{d,BP} - v_{d,SQ} = \zeta^*_d(1-\bar{s}_{-d,BP}) - \zeta^*_d(1-\bar{s}_{-d,SQ})$ for Chrome users. For Edge users, we also need to account for the fact that the utility of using Bing changed by \$10, so $v_{d,BP} - v_{d,SQ} = \zeta^*_d(1-\bar{s}_{-d,BP}) - \zeta^*_d(1-\bar{s}_{-d,SQ}) + \eta\Delta p$.

Data Sharing True qualities are $\zeta^*_G(\bar{s}_G)$ and $\zeta^*_B(1)$. True utilities are $\Delta v_{DS} = \zeta^*_{-d} - \zeta^*_d - \sigma(1-\delta)$. The bias term $\Delta b_{DS} = \tilde{\zeta}_{-d} - \zeta^*_{-d}$ is as in the Status Quo. The true quality of the original search engine changes because of data sharing, so $v_{G,DS} - v_{G,SQ} = \zeta^*_G(\bar{s}_{G,DS}) - \zeta^*_G(\bar{s}_{G,SQ})$ and $v_{B,DS} - v_{B,SQ} = \zeta^*_B(1) - \zeta^*_B(\bar{s}_{B,SQ})$.

Data Sharing + Bing Default + Delayed Choice Screen Everything is as in the Bing Default + Delayed Choice Screen, except that qualities are $\zeta^*_G(\bar{s}_{G,DDD})$ and $\zeta^*_B(1)$.

E.2.1 Computing equilibria

In each of the above counterfactuals, we can plug in the above expressions into our expression for market shares to obtain the following expression for market shares among users of browser b:

$$s_{b,-d} = S(\Delta v_{b,\mathscr{C}}(\bar{s}_{-d}) + \Delta b_{b,\mathscr{C}}(\bar{s}_{-d})).$$

We can aggregate those market shares to obtain Google's total market share

$$\bar{s}_G = \frac{n_{CH} S(\Delta v_{CH,\mathscr{C}}(1-\bar{s}_G) + \Delta b_{CH,\mathscr{C}}(1-\bar{s}_G)) + n_{ED} S(\Delta v_{ED,\mathscr{C}}(\bar{s}_G) + \Delta b_{ED,\mathscr{C}}(\bar{s}_G))}{n_{CH} + n_{ED}},$$

where n_{CH} and n_{ED} represents the number of users on Chrome and Edge, respectively.

Finding an equilibrium consists of computing a solution $\bar{s}_{G,\mathscr{C}}$ to the above equation. We implement this using the bisection method with shares one and zero as starting points. Once we obtain such solution, it is straightforward to compute qualities, which we can use to compute $\Delta v_{\mathscr{C}}$, $\Delta b_{\mathscr{C}}$, and $(v_{d,\mathscr{C}} - v_{d,SQ})$. as well as equilibrium market shares and consumer surplus.

E.3 Additional Counterfactual Simulation Results

Table A6: **Counterfactual Simulations: Strong Effect on Search Result Relevance**

Panel A: Benchmarks						
	Combined		Chrome		Edge	
Description	(1) Google share (%)	(2) CS gain ($/user-year)	(3) Google share (%)	(4) CS gain ($/user-year)	(5) Google share (%)	(6) CS gain ($/user-year)
Status Quo	88.9	0.00	98.8	0.00	22.2	0.00
No Frictions	72.9	6.04	78.8	0.85	33.4	41.04
Active Choice	89.1	5.35	97.3	0.10	33.4	40.67
Correct Perceptions	77.9	0.47	86.2	0.49	22.1	0.35
Correct Perceptions + Data Sharing	77.1	0.72	85.3	0.64	22.1	1.23
No Frictions + Data Sharing	71.9	6.31	77.6	1.06	33.3	41.69
Panel B: Policy Interventions						
	Combined		Chrome		Edge	
Description	(1) Google share (%)	(2) CS gain ($/user-year)	(3) Google share (%)	(4) CS gain ($/user-year)	(5) Google share (%)	(6) CS gain ($/user-year)
Choice Screen	87.5	0.09	97.2	0.10	22.2	0.06
Bing Default	47.7	-70.61	51.5	-81.20	22.1	0.74
Bing Default + Delayed Choice Screen	71.4	0.11	78.7	0.06	22.1	0.48
Bing Payments ($10)	51.3	109.04	56.4	93.95	17.0	210.79
Data Sharing	88.8	0.18	98.8	0.02	22.1	1.26
Data Sharing + Bing Default + Delayed Choice Screen	70.4	0.38	77.6	0.25	22.1	1.21

Notes: This table presents the equilibrium effects of the counterfactual simulation results from the procedure described in Section 7 under the scenario where all consumers experience a strong effect on search result relevance at the fifth percentile given our economy of scale analysis. CS gain means consumer surplus gain, in $/consumer-year. The click-through-rate (CTR) used in the calculation of the equilibrium effects is the average consumer-level click-through-rates (CTR) associated with top organic link clicks.