

Opinion Dynamics via Search Engines (and other Algorithmic Gatekeepers)*

Fabrizio Germano[†] and Francesco Sobbrío[‡]

First Version: December 2016

This Version: April 2017

Abstract

Ranking algorithms are the information gatekeepers of the Internet era. We develop a stylized framework to study the effects of ranking algorithms on opinion dynamics. We consider rankings that depend on popularity and on personalization. We find that popularity driven rankings can enhance asymptotic learning while personalized ones can both inhibit or enhance it, depending on whether individuals have common or private value preferences. We also find that ranking algorithms can contribute towards the diffusion of misinformation (e.g., “fake news”), since lower ex-ante accuracy of content of minority websites can actually increase their overall traffic share.

Keywords: Search Engines, Ranking Algorithm, Search Behavior, Opinion Dynamics, Information Aggregation, Asymptotic Learning, Misinformation, Polarization, Website Traffic, Fake News.

JEL Classification: D83, L86

*We thank Larbi Alaoui, Stefano Colombo, Robert Epstein, Lisa George, Roberto Imbuzeiro Oliveira, Matthieu Manant, Andrea Mattozzi, Ignacio Monzón, Antonio Nicolò, Nicola Persico, Christian Peukert, Emanuele Tarantino, Greg Taylor and seminar participants in Barcelona (UPF), Florence (EUI), Geneva (EEA 2016), Hamburg (Economics of Media Bias Workshop 2015), Naples (Media Economics Workshop 2014), Padova (UP), Palma de Mallorca (JEI 2016), Paris (ICT Conference 2017), Petralia (Applied Economics Workshop 2015) and Rome (IMEBESS 2016) for insightful comments and conversations. We are also indebted to Riccardo Boscolo for helpful discussions on the functioning of search engines. Germano gratefully acknowledges financial support from the Spanish Ministry of Economy and Competitiveness (grant ECO2014-59225-P), from Fundación BBVA (grant “Innovación e Información en la Economía Digital”) as well as from the Barcelona GSE Research Network and the Generalitat de Catalunya.

[†]Department of Economics and Business, Universitat Pompeu Fabra, and Barcelona Graduate School of Economics, fabrizio.germano@upf.edu

[‡]Department of Economics and Finance, LUISS “G. Carli”, fsobbrio@luiss.it

1 Introduction

Search engines are among the most important information gatekeepers of the Internet era. Google alone receives over 3.5 billion search queries per day, and, according to some estimates, 80% of these queries are informational, dwarfing navigational and transactional searches (Jansen *et al.*, 2008). Individuals increasingly use search engines to look for information on a vast array of topics such as science (Horrigan, 2006), birth control and abortion (Kearney and Levine, 2014) or the pros and cons of alternative electoral outcomes (e.g., the Brexit referendum in the UK or the constitutional referendum in Italy; see *Google Trends*). The most remarkable feature of search engines is constituted by their automated algorithm, which determines the ranking of websites to be displayed for any given search query. Similarly, social media platforms—such as Facebook and Twitter—rely on algorithms to rank their items (e.g., posts/tweets). In short, the algorithms used by search engines and social media establish the order of information sources observed by any individual, making them into *de facto* “algorithmic gatekeepers” (Introna and Nissenbaum, 2000; Rieder, 2005; Granka, 2010; Napoli, 2015; Tufekci, 2015).

Despite the widespread presence and relevance of these online platforms and of their gatekeeping algorithms, opinion dynamics via endogenous rankings is largely understudied. This paper aims to fill this gap by developing an empirically grounded, dynamic framework, where individuals use a search engine that ranks websites to obtain information on the state of the world (e.g., side effects of a vaccine), and where the individuals’ choices feed back into the search engine’s rankings, thereby affecting future searches. We then ask how the parameters describing the search environment affect website traffic and asymptotic learning.

The model focuses on three key components of the ranking algorithm: (i) the websites’ initial ranking (reflecting the *ex-ante* parameters that search engines exploit to establish the intrinsic “authority” of websites); (ii) the evolution of websites’ rankings as being dependent on their *popularity*; (iii) the *personalization* of search results according to the individuals’ characteristics (e.g., their geographical location or their web browsing history). At the same time, we make three assumptions on individuals’ online search behavior, namely: (a) they have a *search cost*, which captures the tendency of individuals to give priority of attention to higher ranked results (as in De Cornière and Taylor 2014; Taylor 2013; Burguet *et al.* 2015; Hagi and Jullien 2014); (b) they have a *preference for like-minded news* (as in Mullainathan and Shleifer 2005; Gentzkow and Shapiro 2010); (c) they are *naïve* with respect to the search engine’s algorithm, that is, they do not make any inference from the websites’ rankings *per se*.

We use the model to address basic questions on the effects of ranking algorithms on opinion dynamics: Can they help individuals reach efficient choices? Do they contribute to the diffusion of misinformation (e.g., “fake news”)? Can the personalization of search results hinder asymptotic learning? We identify conditions under which the answer to all three questions is positive.

The popularity component of the algorithm combined with the individual search costs creates a *rich-get-richer* dynamic. Interacted with the individuals’ preference for like-minded news, the *rich-get-richer* dynamic can foster information aggregation. The more individuals have a correct private signal, the more likely they are to visit a website reporting correct information. Consequently, websites with a correct signal will go up in the ranking, and other individuals will be more likely to access a website with a correct signal. Compared to a situation where individuals observe websites ordered randomly, the ranking algorithm enhances the probability of asymptotic learning. Though if the confirmation bias or the preference for like-minded news is very strong, then a random ranking might actually do better than a popularity driven one.

Nevertheless, while ranking algorithms have an overall positive effect on the efficiency of individual choices, at the same time, they can tend to favor websites that are (ex-ante) more likely to report wrong information. In particular, we show that fewer websites carrying a given information may attract more traffic overall, than if there were more of them (*advantage of the fewer*). This implies that, while it is always preferable to have a majority of websites reporting correct information, a larger majority might actually decrease the probability of individuals reading correct information. Indeed, a smaller minority of websites reporting wrong information (e.g., “conspiracy” or “fake news” websites) might each attract a larger share of traffic by like-minded individuals. Remarkably, in the context of algorithmic ranking, this static effect is amplified over time by the *rich-get-richer* dynamic mentioned above. Consequently, the *advantage of the fewer* provides a rationale to explain why “fake news” may thrive and gain in authority in the current information environment, dominated by algorithmic gatekeepers, such as search engines and social media.¹

We also show that the preference for like-minded news combined with the personalization of search results generate *belief polarization*. Most importantly, in a common value model (i.e., when individuals use the search engine to gather information on a state of the world they all equally care about), personalization is likely to inhibit asymptotic learning. Instead, when individuals care differently about different attributes of the state of the world (e.g., different characteristics of a product) an appropriately personalized search algorithm is likely to be efficiency enhancing for individuals compared to a non-personalized one. Therefore, our results suggest that personalization of search results may be sub-optimal in the context of individuals looking for information on common value issues (e.g., side effects of a vaccine), while it may be a useful tool when it comes to search queries on private value issues (e.g., attributes of a commercial product).

We conclude with one important remark. While our model focuses on search engines, our results apply more broadly. The dynamics uncovered by our theoretical results apply to other online platforms using ranking algorithms that act as “algorithmic gatekeepers” such as social

¹Allcott and Gentzkow (2017) document that, in the run-up to the 2016 US presidential election, more than 60% of traffic of fake news websites in the US came from referrals by algorithmic gatekeepers (i.e., search engines and social media).

media platforms (Napoli, 2015; Tufekci, 2015). Indeed, the rankings of items provided by social media depend, among other things, on their popularity. At the same time, the rankings provided are typically personalized, taking user profiles into account. Accordingly, our results suggest that in determining “what is relevant” for their users, these algorithms will also create opinion dynamics featuring patterns of *rich get richer*, *advantage of the fewer* and *belief polarization* as described above.

Related Literature.

To our knowledge, ours is the first paper in economics to analyze opinion dynamics via endogenous algorithmic rankings and the informational gate-keeping role of the search engines.²

The paper is broadly related to the economics literature on the aggregation of information dispersed across various agents (see Acemoglu and Ozdaglar 2011 for a survey). We share with this literature the focus on understanding the conditions under which information that is dispersed among multiple agents might be efficiently aggregated. Our framework differs from this literature in a simple and, yet, crucial aspect: we are interested in investigating the role played by a specific (yet extensively used) “tool” of information diffusion/aggregation, namely, the ranking algorithm, which is used by many online platforms. Most of the literature so far has focused on other mechanisms of information diffusion such as social networks or, more generally, situations where individuals may observe other individuals’ actions (Banerjee, 1992; Bikhchandani *et al.*, 1992; Smith and Sørensen, 2000; Acemoglu *et al.*, 2011; Mueller-Frank, 2013).

Formally, our model of the search engine’s ranking algorithm—and its dynamic evolution—builds on models of individuals’ choices over ranked items used in Demange (2012, 2014). The focus of the paper on the role of search engines as information gatekeepers is close in spirit to the economic literature on news media (see DellaVigna and Gentzkow 2010; Prat and Strömberg 2013, for surveys). At the same time, the presence of an automated ranking algorithm makes search engines—and other algorithmic gatekeepers—fundamentally different from news media. In the case of news media, the choice of what information to gather and disclose is made on a discretionary, case-by-case basis. In the case of search engines the gate-keeping is unavoidably the result of an automated algorithm (Granka, 2010; Tufekci, 2015).³ Therefore, whatever *bias* might originate from search engines, its nature is intrinsically different from the one arising from news media, especially in traditional media. At the same time, this implies that studying the effects of search engines on the accuracy of individuals’ beliefs requires a rather different approach with respect

²The existing economics literature on search engines has focused on the important case where search engines have an incentive to distort sponsored and organic search results in order to gain extra profits from advertising and product markets (Taylor, 2013; De Cornière and Taylor, 2014; Hagiu and Jullien, 2014; Burguet *et al.*, 2015). See also Grimmelmann (2009) and Hazan (2013) for a legal perspective on the issue.

³Put differently, “While humans are certainly responsible for editorial decisions, these [search engine] decisions are mainly expressed in the form of software which thoroughly transforms the ways in which procedures are imagined, discussed, implemented and managed. In a sense, we are closer to *statistics* than to *journalism* when it comes to bias in Web search”; Rieder and Sire (2013), p. 2.

to the one present in theoretical models of media bias, (e.g., Strömberg, 2004; Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006).

Finally, the paper is also related to the literature outside economics discussing the possible implications of the search engines' architecture (and, more generally, of algorithmic gatekeeping) on democratic outcomes. This literature encompasses communication scholars (Hargittai, 2004; Granka, 2010), legal scholars (Goldman, 2006; Grimmelmann, 2009; Sunstein, 2009), media activists (Pariser, 2011), psychologists (Epstein and Robertson, 2015), political scientists (Putnam, 2001; Hindman, 2009; Lazer, 2015), sociologists (Tufekci, 2015) and, last but not least, computer scientists (Cho *et al.*, 2005; Menczer *et al.*, 2006; Glick *et al.*, 2011; Pan *et al.*, 2007; Yom-Tov *et al.*, 2013; Flaxman *et al.*, 2013; Bakshy *et al.*, 2015).

The paper is structured as follows. Section 2 describes the framework. Section 3 provides a background discussion of the main assumptions of the model. Section 4 presents the main results concerning the effects of the search engine's algorithm on the evolution of the website ranking and on individuals' opinion dynamics, for the case where there is no personalization of search results. Section 5 discusses the effect of personalization of search results on belief polarization. Section 6 presents the implications of the model in terms of asymptotic learning, also providing a comparison with a fully randomized ranking. Section 7 introduces some extensions which assess the robustness of the results of the benchmark model and provide further insights, and Section 8 concludes. All the proofs and the some formal definitions are relegated to the Appendix.

2 The Model

We present a stylized model of a search environment where individuals use a search engine to seek information on a fixed issue. Individuals all have the same preference and perform the same query to choose an action that matches the true state of the world. This action may have several interpretations. For example, it may represent an individual's choice of whether to vaccinate her child against measles, whether to vote in favor of a referendum or electoral candidate and so on. Individuals get their information from websites, which simply report their own private signal, assumed to be constant throughout. At the center of the model is a search engine that is characterized by its ranking algorithm, which ranks and thereby directs individuals to different websites, using, among other things, the popularity of what individuals previously searched and accessed. To simplify the analysis, we assume that individuals perform exactly one search, one after the other. After that they choose to access one website, whose information they use to update their beliefs and make a decision about their optimal action. We now describe the formal environment.

2.1 Information Structure

There is a binary state of the world ω , which is a $(\frac{1}{2}, \frac{1}{2})$ Bernoulli random variable which takes one of two values from the set $\{0, 1\}$. There are M websites and N individuals, where, by slight abuse of notation, we let $M = \{1, \dots, M\}$, $N = \{1, \dots, N\}$ also denote the set of websites and individuals, respectively. Each website $m \in M$ receives a private random signal correlated with ω ,

$$y_m \in \{0, 1\} \text{ with } \mathbb{P}(y_m = \omega \mid \omega = z) = q \in (\frac{1}{2}, 1), \text{ for any } z \in \{0, 1\}.$$

Similarly, each individual $n \in N$ receives a private random signal,

$$x_n \in \{0, 1\} \text{ with } \mathbb{P}(x_n = \omega \mid \omega = z) = p \in (\frac{1}{2}, 1), \text{ for any } z \in \{0, 1\}.$$

We assume signals to be independent across individuals and websites when conditioned on ω , and assume, moreover, $q > p$, reflecting the fact that, from an ex-ante perspective, the signal received by any given website is more accurate than the one received by any given individual.

2.2 Websites

Each of the $M > 1$ websites is characterized by a signal $y_m \in \{0, 1\}$. Websites are assumed to report their signal and to keep it fixed throughout the search process. They are therefore nonstrategic players that carry information that they provide for free to any individual accessing their website.⁴ For simplicity and unless otherwise noted, we assume that M is an odd number.

2.3 Individuals

Each of the N individuals must choose a binary action $a_n \in \{0, 1\}$, which gives a payoff of one, if the action matches the true state ($a_n = \omega$), and gives a payoff of zero otherwise,

$$u_n(a_n, \omega) = \begin{cases} 1 & \text{if } a_n = \omega \\ 0 & \text{else.} \end{cases}$$

They do not know the true state of the world ω , and have to make their decision based on their information, which consists of their private signal x_n and the signal y_m of the website they access through the search engine. We assume individuals perform their searches once, sequentially, and enter in a random *order* such that, at any point in time t , there is a unique individual $n \equiv t$, who receives a random (i.i.d.) signal x_n as specified above.⁵ As explained in Appendix A this

⁴Besides allowing for more tractability as it avoids modeling the strategic decisions of choosing which signal to report, the assumption also reflects the fact that many websites can represent an article or a document that is already posted and that contains pertinent information to the given search query. Our analysis focuses on which of the different websites or information sources are ultimately accessed by individuals.

⁵For reasons of tractability, we assume that all individuals perform the same search query exactly once.

assumption allows us to study the limit properties of opinion dynamics via search engines from an *ex-ante* perspective, by approximating the stochastic search process using the *mean dynamics* methodology.⁶

2.3.1 Ranking-free website choice

Consider an individual who has no search cost and who needs to choose which website to consult. With a zero search cost, the websites ranking does not matter for the individual and her website choice is what we call *ranking-free*. If, moreover, the individual has no preference for like-minded news, then her best option is to simply consult a website that carries the signal that has the highest ex ante probability of being correct. With only two possible signals, and given $q > p$, the signal carried by the majority of websites, the *website majority signal*, denoted by $y_K \in \{0, 1\}$, plays this role. Let $K = \{m \in M \mid y_m = y_K\}$, where $\frac{M}{2} < K \leq M$, denote the set (and number) of websites carrying the website-majority signal. Then, a rational individual would want to consult a website in K regardless of her own signal. If, on the other hand, the individual has a preference for like-minded news, then she may want to consult a website that carries the same signal as her own even when this signal is not the website-majority one. In particular, we assume that the *preference for like-minded news* is captured by a parameter $\gamma \in [0, 1]$, which represents the probability of an individual choosing a website that carries the same signal as her own, regardless of whether this signal is a website-majority signal or not. The following table describes the probabilities that an individual with signal x_n will choose given website with signal y_m , when $K \subsetneq M$.

Table 1: Ranking-free choice probabilities $\rho_{n,m}^*$

		Website signal	
		$y_m = y_K$	$y_m = y_{M \setminus K}$
Individual signal	$x_n = y_K$	$\frac{1}{K}$	0
	$x_n = y_{M \setminus K}$	$\frac{1-\gamma}{K}$	$\frac{\gamma}{M-K}$

Thus, if an individual's private signal coincides with the website-majority signal ($x_n = y_K$), then she clicks on any website in K with probability $\frac{1}{K}$. If it does not coincide with the website-majority signal ($x_n \neq y_K$), then she clicks on any website in K only with probability $\frac{1-\gamma}{K}$. With this we can define the *ranking-free choice function* $\rho_n^* : \{0, 1\} \rightarrow \Delta(M)$,⁷

$$\rho_{n,m}^*(x_n) = \frac{\gamma \mathbb{I}_{\{y_m = x_n\}} + (1 - \gamma) \mathbb{I}_{\{y_m = y_K\}}}{\sum_{m'} \mathbb{I}_{\{y_m = y'_m\}}}, \quad (1)$$

⁶Clearly, whenever individuals do not enter randomly, opinion dynamics may further depend on the specific order of the individual searches.

⁷To avoid excessive notation, we omit the arguments y_m and y_K from the function ρ_n^* .

which, for any $m \in M$, and in the absence of search costs, gives the ex-ante probability of individual n choosing website m .

2.3.2 Search cost

When performing their search, agents obtain a list of *ranked* websites, which they need to process at a cost. Thus, besides their *ranking-free* preferences over websites, contingent on a website signal and on their own private signal, if individuals have a search cost, they may implicitly favor some websites over others, depending on those websites' position in the observed ranking. Let $r_t \in \Delta(M)$ be the *ranking* of the M websites at time t (i.e., faced by individual $n \equiv t$). The ranking is provided by the search engine (as discussed in the next section). When processing the ranked websites, we assume the individuals use the following *weighting function* $\sigma : \Delta(M) \rightarrow \Delta(M)$, $\sigma(r_n) = (\sigma_1(r_n), \dots, \sigma_M(r_n))$, where, for $m \in M$,

$$\sigma_m(r_n) = \frac{(r_{n,m})^\alpha}{\sum_{m'} (r_{n,m'})^\alpha}, \text{ for some } \alpha \geq 0. \quad (2)$$

The parameter α calibrates the individual's *search cost* in the following sense: $\alpha > 1$ magnifies the difference in the entries of r_t and therefore magnifies the differences present in the initial ranking r_1 ; $\alpha < 1$ reduces the difference in the entries of r_t and hence in the initial ranking r_1 ; in the limit as $\alpha \rightarrow 0$ all entries have the same weight, which represents the case with *no* search cost, where all websites that provide the same signal are accessed with the same *ex ante* probability; finally, $\alpha = 1$ maintains the original differences present in r_t .

2.3.3 Website choice

We are finally in a position to define the actual website choices of individuals. This is embodied in the *website choice function* $\rho_n : \{0, 1\} \times \Delta(M) \rightarrow \Delta(M)$, defined by,⁸

$$\rho_{n,m}(x_n, r_n) = \frac{\sigma_m(r_n) \cdot \rho_{n,m}^*(x_n)}{\sum_{m'} \sigma_{m'}(r_n) \cdot \rho_{n,m'}^*(x_n)} = \frac{(r_{n,m})^\alpha \cdot \rho_{n,m}^*(x_n)}{\sum_{m'} (r_{n,m'})^\alpha \cdot \rho_{n,m'}^*(x_n)}, \quad (3)$$

which, for any $m \in M$, gives the probability that individual n accesses website m as a result of his search at time $t \equiv n$. It is obtained from the ranking-free choice function ρ_n^* weighted by the function σ , normalized to give a probability distribution on the set of websites M .

Essentially, the *weighting function*—and the associated search cost— augments or reduces the ranking differences between different websites contained in r_t , thus affecting the way the probabilities of visiting websites are distributed among all ranked websites (e.g., when $\alpha = 0$, we have $\rho_n \equiv \rho_n^*$). The individual preference over websites—reflected in the *ranking-free* choice function— on the other hand can act to eliminate websites that are not visited; or reduce the probability

⁸Fortunato *et al.* (2006) and Demange (2012, 2014) use related models of individuals' choices over ranked items.

with which they are visited. In particular, for $\gamma = 0$, ρ_n^* represents the incentive-compatible choice of individual n : with probability one she visits one of the websites carrying the website-majority signal. Hence, absent biases, individuals consistently use the distribution of the websites' signals to make their choices.⁹ When including a preference for like-minded news ($\gamma > 0$), a behavioral bias is introduced such that ρ_n^* may be seen as satisfying an incentive compatibility requirement distorted by a behavioral bias, namely the preference for like-minded news (Rabin and Schrag, 1999; Mullainathan and Shleifer, 2005).¹⁰

2.4 Search Engine and Ranking Algorithm

The *ranking algorithm* \mathcal{R} used by the search engine is at the center of our model. At each point in time, $t = 1, 2, \dots, N$, it provides a *ranking* $r_t \in \Delta(M)$ of the M websites, where an element $r_{t,m}$ is the probability that the individual searching at time t is directed to website m in the absence of any other factors. Concretely, we assume that given an initial ranking $r_1 \in \Delta(M)$, the ranking r_t at subsequent periods, $t = 2, 3, \dots, N$, is defined by,

$$r_{t,m} = \nu r_{t-1,m} + (1 - \nu) \rho_{t-1,m}, \quad (4)$$

where $\nu \in (0, 1)$. This says that the ranking at time t , r_t , is determined by the ranking of the previous period r_{t-1} and partly also by the individual's choice over websites in the previous period $\rho_{t-1,m}$ as defined in (3).¹¹ The weight that is put on each term depends on the parameter ν , where for convenience we write,

$$(\nu, 1 - \nu) = \left(\frac{\kappa}{\kappa + 1}, \frac{1}{\kappa + 1} \right), \quad (5)$$

where in turn $\kappa \in \mathbb{N}$ is a *persistence* parameter of the ranking algorithm \mathcal{R} . The higher κ , the more persistent and stable the search engine's ranking is. In Section 6, when studying asymptotic behavior (as $N \rightarrow \infty$), we will assume $\kappa \rightarrow \infty$.

⁹This can be seen to lead to limiting behavior consistent with an analogy-based expectations equilibrium along the lines of Guarino and Jehiel (2013).

¹⁰More generally, an element $\rho_{n,m}^*$ of $\rho_n^*(x_n) \in \Delta(M)$ is the probability that individual n would want to visit website m knowing her own private signal x_n and the signals $(y_m)_{m \in M'}$ of all websites in a subset $M' \subset M$, and given the preference for like-minded news γ . The idea is that besides her own signal, an individual observes the headlines (or *snippets*) of a number of websites, M' , both of which are reflected in the *ranking-free website choice function* $\rho_{n,m}^*$, such that when not following her own signal she chooses a website carrying the signal most represented among websites. For simplicity, we assume that $M' \equiv M$ consists of all websites and discuss the more realistic case where the website choice is based on a smaller subset of websites $M' \subset M$, from which individuals “perceive” which is the website-majority signal, is discussed in Section 7. The latter can be seen as introducing an additional behavioral bias in individual preferences (Tversky and Kahneman, 1974).

¹¹This updating algorithm via the “popularity” of a website maybe interpreted both in a strict sense (e.g., direct effect of the actual clicks on the website in the search result page) and in a broad sense (e.g., a website that receives more clicks is also more likely to be more popular in other online platforms and viceversa).

2.5 Personalized Ranking Algorithm

An important question for a ranking algorithm concerns whether it should keep track of, and use, information it has available concerning the identity of individuals performing searches. A search engine may want its algorithm to condition the outcomes of searches on the geographical location of the individuals (e.g., using the individual’s *IP address*) or on the individual’s socio-economic profile (e.g., using the individual’s search history). Accordingly, a personalized search algorithm may output different search results to the same query performed by individuals living in different locations and/or with different browsing histories.¹²

Suppose the set of individuals N is partitioned into two nonempty groups $A, B \subset N$, such that $A \cup B = N$ and $A \cap B = \emptyset$. In any period an individual is randomly drawn from one of the two groups, that is, from A with probability $\frac{N_A}{N}$ and from B with probability $\frac{N_B}{N}$, where $N_A = \#A$ and $N_B = \#B$. A *personalized ranking algorithm* \mathcal{R}_ℓ then consists of two parallel rankings, namely, r_t^A for individuals in A and r_t^B for individuals in B . Each one is updated as in the non-personalized case, with the difference that the weight on past choices of individuals from the own group are possibly larger than those from the other one. Set, for any t and m , and for $\ell = A, B$,

$$r_{t,m}^\ell = \nu_t^\ell r_{t-1,m} + (1 - \nu_t^\ell) \rho_{t-1,m}^\ell, \quad (6)$$

where ν_t^ℓ now depends on whether or not the individual searching at time $t - 1$ was in the same group, $t - 1 \in \ell$, and where the weight ν_t^ℓ , for $\ell = A, B$, is given by:

$$(\nu_t^\ell, 1 - \nu_t^\ell) = \left(\frac{\kappa}{1 - \lambda_t + \kappa}, \frac{1 - \lambda_t}{1 - \lambda_t + \kappa} \right), \text{ where } \lambda_t = \begin{cases} 0 & \text{if } t - 1 \in \ell \\ \lambda & \text{else,} \end{cases} \quad (7)$$

with $\kappa \in \mathbb{N}$ and $\lambda \in [0, 1]$ parameters of the personalized ranking algorithm \mathcal{R}_ℓ . This algorithm now gives different weights to past choices over websites depending on whether these choices were taken by individuals in the same group (weights $\frac{\kappa}{1+\kappa}, \frac{1}{1+\kappa}$) or in the other group (weights $\frac{\kappa}{1-\lambda+\kappa}, \frac{1-\lambda}{1-\lambda+\kappa}$). Thus, when $\lambda = 1$, the ranking algorithm is fully personalized, whereas, when $\lambda = 0$, it coincides with the non-personalized one previously defined. We implicitly assume that the personalized search algorithm partially separates individuals according to some individual characteristics. In particular, we will consider the case, where individuals are distinguished based on the accuracy of their private signals (i.e., $p^A \neq p^B$).

2.6 Search Environments

A search environment is to be thought of as an *ex ante* notion that fixes the ranking algorithm, information structure and characteristics of individuals and websites, before they receive their

¹²See Pariser (2011); Dean (2013); MOZ (2013); Vaughn (2014); Hannak *et al.* (2013); Xing *et al.* (2014); Kliman-Silver *et al.* (2015).

signals and before they perform their search. More formally, we define a *search environment* \mathcal{E} as a list of variables,

$$\mathcal{E} = \langle (p, q); (N, \alpha, \gamma); M; (\kappa, \lambda) \rangle, \quad (8)$$

where (p, q) describes the information structure, (N, α, γ) describes the individuals, M describes the websites, and (κ, λ) describes the ranking algorithm. Unless otherwise noted, we will use the term search environment to denote the case with a non-personalized ranking algorithm ($\lambda = 0$ or ranking described by Equation (4)) and use the term *personalized search environment* for the case where the weights depend on the group ($\lambda > 0$ and the ranking described by Equation (6)). It will be clear from the context when we are considering environments with personalized rankings. Given a search environment \mathcal{E} , whether personalized or not, we refer to an (*interim*) *realization* of \mathcal{E} as to the tuple $\langle \omega; (L, (y_m)_{m \in M}); (x_n)_{n \in N} \rangle$, where the true state of the world and the signals of the websites are fixed; L denotes the number of websites with the correct signal $y_m = \omega$; individuals with signals x_n enter sequentially, one at a time, in a random order.

We let r_1 denote the *initial ranking*. Unless otherwise specified, we assume that r_1 is interior, that is, $r_{1,m} > 0$ for all $m \in M$. Let also $r_{t,K}$ and $r_{t,M \setminus K}$ denote the total probabilities put by the ranking algorithm at time t on all websites reporting the website-majority and the website-minority signal, respectively; similarly for the clicking probabilities $\rho_{t,K}$ and $\rho_{t,M \setminus K}$. We also often talk about the *expected* probability of individual n accessing website m , $\hat{\rho}_{t,m} = \mathbb{E}[\rho_{t,m}]$, where the expectation is taken over the private signal of the agent n that enters to perform a search at time $t \equiv n$. These probabilities are discussed more formally in Appendix A (see, in particular, Equation (14)).

3 Background and Interpretation of the Model

In this section we motivate and discuss the key assumptions of our model.

3.1 Individual Search Behavior.

Our objective is to study the dynamic interaction of ranking algorithm and individual informational searches.¹³ The modeling strategy is to work with a tractable, parameterized and reduced form model of naïve or boundedly rational choice on the part of the individuals and a constant strategy on the part of the websites. In particular, the model makes three key assumptions on the search behavior of individuals: (a) individuals incur a search cost when looking for information; (b) they have preference for like-minded news; and (c) they are naïve with respect to the observed website rankings. Assumption (a) does not constitute a deviation from a rational-choice framework, unlike (b), (c), which introduce behavioral elements. However, while the *naïveté* of individuals with respect

¹³Our analysis is positive and we do not pursue a mechanism design approach here. Kremer *et al.* (2014) solve for an optimal disclosure policy in the context of a dynamic recommendation system. Palacios-Huerta and Volij (2004) and Altman and Tennenholtz (2008) study axiomatizations of static ranking systems.

to the ranking is made to simplify the analysis, the assumption on the preference for like-minded news allows to obtain further insights on opinion dynamics with ranking algorithms. We discuss these assumptions in more detail.

As in the literature on search diversion (Taylor 2013; De Cornière and Taylor 2014; Hagiü and Jullien 2014; Burguet *et al.* 2015), individuals have a *search cost* (captured by the parameter α), which, all other things equal, increases their probability of choosing a higher ranked website. There is plenty of evidence showing that a large majority of individuals tend to disregard search results beyond the first page (Spink and Jansen, 2004; Guan and Cutrell, 2007; Chitika, 2013). As shown by Pan *et al.* (2007), Glick *et al.* (2011) Yom-Tov *et al.* (2013) and Epstein and Robertson (2015), this is not just a mere correlation. Keeping all other things equal (e.g., the fit of a given website with respect to the individual’s preferences), highest ranked results tend to receive significantly more “attention” by users than lower ranked ones.¹⁴

We also allow individuals to have a *preference for like-minded news* (represented by the parameter γ). This captures a further important behavioral aspect of how individuals choose among given search results. Online Appendix B shows that this reduced-form parametrization is open to alternative interpretations. For example, the preference for like-minded news may result from a disutility from reading news inconsistent with the individuals’ beliefs (as in Mullainathan and Shleifer 2005) or, alternatively, they may be the result of a confirmation bias (as in Rabin and Schrag 1999). This assumption is also justified on empirical grounds. Gentzkow and Shapiro (2010) document that readers of US newspapers have an economically significant preference for like-minded news.¹⁵

Individuals are assumed to be naïve: they do not make inferences on the informativeness of a website’s signal based on the observed ranking of that website. Accordingly, our theoretical framework does not encompass a fully-fledged model of rational herding (Banerjee, 1992; Bikhchandani *et al.*, 1992; Smith and Sørensen, 2000). Shutting down the possibility of individuals making inferences via the websites’ ranking *per se* is clearly a simplifying assumption in our model. While we think that this assumption is not far-fetched, we are not aware of any empirical evidence supporting or disproving it. In this sense, our model follows the literature that studies learning processes with information processing biases (DeMarzo *et al.*, 2003; Eyster and Rabin, 2010; Guarino and Jehiel, 2013).¹⁶ At the same time, even though individuals are naïve, the interaction between the individ-

¹⁴Epstein and Robertson (2015) study a related phenomenon they refer to as the “search engine manipulation effect” (SEME). In a series of experiments on voters’ opinions and preferences, they show how changing the ranking of websites may have a potentially very large impact on political outcomes. More generally, recent research shows that, when faced with ordinal lists, individuals often show a disproportionate tendency to select options that are placed at the top. See Novarese and Wilson (2013) for a discussion of the literature.

¹⁵See also Yom-Tov *et al.* (2013); Flaxman *et al.* (2013); White and Horvitz (2015) for similar evidence in the specific context of search engines.

¹⁶Several models of social or observational learning have been proposed that depart from the original fully rational Bayesian case. Monzón and Rapp (2014) study a model of rational Bayesian learning where individuals sample the decisions of past individuals but have uncertainty or only partial information about their position in the sequence of choices. Bohren (2016) studies learning by individuals with misspecified models. Eyster and Rabin (2010) study agents who believe previous agents’ actions only reflect their signals. Guarino and Jehiel (2013) study analogy-

ual search cost and the “popularity” component of the ranking algorithm still induces information aggregation in the model: the more people receive a correct signal, the more likely other people are to choose a website reporting such a signal (see Section 6). Online Appendix A briefly discusses the case of sophisticated individuals.

Finally, we remark that, while individuals observe whether a website is reporting a “website-majority” or a “website-minority” signal, they still need to click on a website to update their beliefs on the issue at stake. This mechanism is in line with the literature on the demand for slanted news (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006; Chan and Suen, 2008; Sobbrío, 2014; Alaoui and Germano, 2016) where individuals know the ideological slant of each source of news but they need to actually watch the news report in order to update their beliefs (e.g., individuals know that Fox News is not favorable to the Obamacare, but they will update their beliefs on the consequences of the reform only upon actually watching the news reports).

3.2 The Search Engine’s Algorithm.

There are two main challenges in constructing a formal model providing a stylized representation of a search engine’s algorithm. First, search algorithms are complex: Google currently uses around 200 signals in determining the ranking of search results for a given query (Dean, 2013; MOZ, 2013; Vaughn, 2014).¹⁷ Second, to prevent the possible abusive use of the algorithms, their exact features typically represent a commercial secret. Nevertheless, while the exact details on these ranking algorithms are kept secret, certain key features constitute “common wisdom” among folk-tech experts. Most importantly, the main components of search algorithms are well known and studied by computer scientists. Many of the factors used by search engines to decide the ranking of a webpage (for a given query) are based on *ex-ante* parameters that the algorithm uses to establish the initial “authority” of that webpage.¹⁸

In order to simplify and preserve tractability, we consider these parameters to be static components of such an algorithm. Accordingly, the initial ranking is implicitly assumed to capture the *ex-ante* ranking of websites for a given search query. A second set of parameters of search engines’ algorithms is instead referring to user-interaction (e.g., click-through-rate) and social signals (e.g., webpage likes on Facebook, tweets linking to the webpage). That is, as specified in Equation (4), search engines update the initial ranking of a webpage according to how “popular” the webpage is. Finally, as captured in Equation (6), search engines’ algorithms personalize search results according

based expectations equilibria, where agents make inferences from the aggregate distribution of choices made without necessarily keeping track of all past history. We believe these are worth exploring further in a context of individuals making choices after observing an endogenous algorithmic ranking.

¹⁷As noted by Lazer (2015) such complexity is amplified by the interaction of the algorithm itself with individual search behavior. In particular, “the interplay of social algorithms and behaviors yields patterns that are fundamentally emergent. These patterns cannot be gleaned from reading code.” (Lazer 2015, p. 1090). In turn, this also suggests a theoretical model studying the interplay between users’ search behavior and the algorithm’s rules may yield insights that may otherwise be difficult to identify simply looking at the algorithm code or at observational data.

¹⁸The *PageRank* algorithm (Brin and Page, 1998) is a particularly important example.

to individual characteristics such as, for example, her geographical location (e.g., IP address) and her past search and browsing behavior (Pariser 2011; Dean 2013; Epstein 2013; MOZ 2013; Vaughn 2014; Hannak *et al.* 2013; Xing *et al.* 2014; Kliman-Silver *et al.* 2015).¹⁹

4 Rich Get Richer, Concentration at the Top, Long-tail and the Advantage of the Fewer

We here analyze the interaction of the ranking algorithm with costly search ($\alpha > 0$) and preference for like-minded news ($\gamma > 0$) in a setting where there is no personalization ($\lambda = 0$). In Section 5 we study personalized rankings ($\lambda > 0$).

4.1 Initial Ranking and the Rich Get Richer

The evolution of a website’s ranking based on its “popularity,” interacted with a sufficiently large search cost ($\alpha > 1$) of the individuals exhibits a *rich-get-richer* dynamic, whereby the ratio of the expected clicking probabilities of two websites m, m' , with $r_{1,m} > r_{1,m'} > 0$ increases over time as more agents perform their search. The effect is further magnified, the larger α is. Importantly, the differences in ranking and, then, in the *expected* probability of accessing a given website, are driven by the initial ranking (r_1) and are amplified by the search cost. We define this more formally.

Definition 1. *Fix a search environment \mathcal{E} with interim realization $\langle \omega; (L, (y_m)) \rangle$. We say that \mathcal{E} exhibits the rich-get-richer dynamic if, for two websites $m, m' \in M$ with the same signal, $y_m = y_{m'}$, and different initial ranking, $r_{1,m} > r_{1,m'} > 0$, we also have $\frac{\hat{\rho}_{n,m}}{\hat{\rho}_{n,m'}} > \frac{\hat{\rho}_{n-1,m}}{\hat{\rho}_{n-1,m'}}$, for any $n > 0$.*

When an environment exhibits a *rich-get-richer* dynamic, then the ratio of the *expected* probability of two websites (with the same signal but different initial ranking) being visited not only persists over time (this follows from $\kappa > 0$), but actually increases. We can state the following.

Proposition 1. *Let \mathcal{E} be a search environment with positive search cost ($\alpha \geq 0$) and with arbitrary preference for like-minded news ($\gamma \geq 0$) and with interim realization $\langle \omega; (L, (y_m)) \rangle$. Then we have, for any two websites $m, m' \in M$ with $y_m = y_{m'}$ and $r_{1,m} > r_{1,m'} > 0$, and $n > 1$,*

$$\frac{\hat{\rho}_{n,m}}{\hat{\rho}_{n,m'}} \begin{cases} < \\ = \\ > \end{cases} \frac{\hat{\rho}_{n-1,m}}{\hat{\rho}_{n-1,m'}} \begin{cases} \text{if } \alpha < 1 \\ \text{if } \alpha = 1 \\ \text{if } \alpha > 1. \end{cases} \quad (9)$$

In particular, if the search cost is large enough ($\alpha > 1$), then \mathcal{E} exhibits the rich-get-richer dynamic.

¹⁹Hannak *et al.* (2013) document the presence of extensive personalization of search results. In particular, while they show that the extent of search results personalization varies across topics, they also point out that “politics” is the most personalized query category. See also Xing *et al.* (2014) for empirical evidence on search results personalization based on the *Booble* extension of Chrome.

This proposition shows that the search cost plays a crucial role in the evolution of website traffic. When $0 \leq \alpha < 1$, initial conditions do not matter in the limit (as $N \rightarrow \infty$), in the sense that in the limit websites with the same signal tend to be visited with the same probability. The case $\alpha = 1$ is a borderline case, where the ratios of the expected clicking probabilities remain constant for websites with the same signal. When $\alpha > 1$, initial conditions matter and the evolution of website traffic follows a rich-get-richer dynamic. Traffic concentrates on the websites that are top ranked in the initial ranking.²⁰

Notice that the *rich-get-richer* dynamic also creates an implicit form of herding in our model. If a website becomes more popular, it goes up in the ranking and, with a high enough search cost, it will also be more likely to be chosen later on. In turn, this implies that if more people receive a majority (minority) signal, other individuals will also have a higher probability of choosing a website reporting a majority (minority) signal. As we will discuss in Section 6, this implies that the popularity component in the ranking algorithm, together with the search cost, can generate information aggregation in the sense that the more people receive a correct signal, the more likely subsequent people are to choose a website reporting such a signal.

The *rich-get-richer* dynamic is in line with the “Googlearchy” suggested by Hindman (2009), who argues that the dominance of popular websites via search engines is likely to be self-perpetuating. Most importantly, the *rich get richer* pattern of website ranking (and traffic) via search engines is consistent with established empirical evidence (Cho and Roy, 2004).²¹

4.2 Concentration at the Top and the Long Tail

The *rich-get-richer* dynamic implies a tendency towards concentration of website traffic at the top. At the same time, the presence of a preference for like-minded news ($\gamma > 0$) may create a *long-tail* in the pattern of website traffic. The following definition provides a formal notion of *concentration at the top* and *long-tail* as properties of a search environment \mathcal{E} .

Definition 2. Fix a search environment \mathcal{E} with realization $\langle \omega; (L, (y_m)) \rangle$ and a non-empty subset of websites $M' \subset M$ and constant $\frac{1}{2} < \zeta < 1$. We say that \mathcal{E} exhibits concentration at the top with M' and ζ , if $\sum_n \rho_{n,M'} > \zeta$. We say \mathcal{E} exhibits a long-tail with M' and ζ , if $\sum_n \rho_{n,M \setminus M'} < 1 - \zeta$ and $\rho_{n,m} > 0$, for any $m \in M \setminus M'$.

Given Proposition 1, it is easy to see how *concentration at the top* and *long-tail* are not only potential properties of search environments, but can actually easily be generated with a sufficiently large search cost ($\alpha > 1$) and a positive preference for like-minded news ($\gamma > 0$), when the initial

²⁰Online Appendix C provides a numerical and graphical illustration of this and of other results of the paper.

²¹Indeed, even if some scholars have argued that the *overall* traffic induced by search engines is less concentrated than it might appear due to the topical content of user queries (Fortunato *et al.*, 2006), the *rich-get-richer* dynamic is still present within a specific topic.

ranking is *generic*, that is, r_1 is such that $r_{1,m} > 0$, for any $m \in M$, and $r_{1,m} \neq r_{1,m'}$ for any two websites $m \neq m'$.

Corollary 1. *Fix a search environment \mathcal{E} with realization $\langle \omega; (L, (y_m)) \rangle$ with generic initial ranking r_1 . Let $1 < k < M$ and $\frac{1}{2} < \zeta < 1$ be given. Then, for N sufficiently large, $\alpha > 1$ and $\gamma > 0$ there will be concentration at the top and a long-tail for some $M' \subset M$ with $M' = k$ and ζ . This is true even for generic initial rankings that are arbitrarily close to being uniform.*

Empirical studies suggest that website traffic is often characterized by a high concentration at the top and a *long-tail* (Hindman 2009; Gentzkow and Shapiro 2011). Our framework identifies two possible forces driving such a pattern of online traffic. On one hand, the interaction of the ranking algorithm with the individuals' search cost can lead to high concentration of traffic on few top websites. On the other, the presence of a preference for like-minded news further contributes towards generating a long-tail.

4.3 The Advantage of the Fewer

We now show a rather general phenomenon within our framework, whereby a set of websites with the same signal can get a *greater* total clicking probability if the set contains fewer websites than if it contains more of them (as long as it does not switch from being a set of majority to a set of minority websites). We refer to this as *advantage of the fewer*.

Definition 3. *Let \mathcal{E} be a search environment with two realizations that only differ in the set of websites carrying the majority signal $K, K' \subset M$, where $\frac{M}{2} < \#K \leq \#K' < M$. Then we say that \mathcal{E} exhibits an advantage of the fewer if $\hat{\rho}_{N,K} \geq \hat{\rho}_{N,K'}$ and $\hat{\rho}_{N,M \setminus K'} \geq \hat{\rho}_{N,M \setminus K}$.*

The idea is that fewer websites carrying the same signal can obtain overall more traffic than if there were more. The property holds under fairly general conditions.

Proposition 2. *Let \mathcal{E} be a search environment with N and κ large, $0 < \gamma < 1$, and with a uniform initial ranking r_1 , then \mathcal{E} exhibits an advantage of the fewer.*

This suggests that in the context of opinion dynamics via search engines, where the ranking of websites matters (e.g., when $\alpha > 0$), having fewer websites reporting a given information may actually enhance the overall ranking and traffic of such websites. In turn, this implies that having a limited majority (or a limited minority) of websites reporting a given signal actually increases their overall traffic since the individual websites benefit from relatively higher rankings.²² At the same

²²One might expect that allowing for free entry would tend to weaken the effect. However, while addressing free entry of websites and even modeling the strategic choices of websites is outside the scope of this paper, we think it is worth noting that the effect of allowing free entry on the *advantage of the fewer* may be limited, especially in those cases where the “fewer websites” are minority websites carrying “dubious” information. Indeed, it may not be in the interest of mainstream/majority websites (which are likely to also care about their reputation) to report such information. At the same time, new minority websites would not be able to divert much traffic from the existing ones due to their lower ranking and the *rich-get-richer* dynamic. Hence we believe the result may be particularly relevant for websites carrying dubious information that “resonates” with a significant fraction of individuals.

time, the amount of traffic that such websites can attract is limited by the fact that majority (e.g., mainstream) websites attract all traffic from individuals with a majority signal. Accordingly, the results do not imply that all traffic is directed toward a single website, but rather that a significant share of traffic can be concentrated on both majority and minority websites. Also, the fewer the websites representing those majority or minority websites, the larger their share of traffic will be.

Given that minority websites are the ones ex-ante more likely to report incorrect information, the *advantage of the fewer* suggests that ranking algorithms may indirectly favor the spread of misinformation. This is consistent with various cases where the algorithms used by Google and Facebook have apparently promoted websites reporting “fake news”.²³ Indeed, even in queries with a clear factual truth, such as yes-no questions within the medical domain, top-ranked results of search engines provide a correct answer less than half of the time (White, 2013). Most importantly, empirical evidence shows that websites reporting misinformation may acquire a large relevance in terms of online traffic and in turn may affect individuals beliefs and behavior (Carvalho *et al.*, 2011; Kata, 2012; Mocanu *et al.*, 2015; Shao *et al.*, 2016).

5 Personalized Ranking Algorithm and Belief Polarization

By introducing personalization in our model, we allow the ranking of websites to be conditioned on (observable) characteristics of the individuals such that searches performed by individuals in different groups can have different weights. When λ is close to zero, there is little difference in the choice over websites of the two groups, while as λ increases the groups may start to have an increasingly different probability of choosing the same type of websites and, thus, to form different opinions. In other words, increased personalization may lead to increased *belief polarization*.

Definition 4. Fix a search environment \mathcal{E} with groups A and B with equal numbers of individuals. Let K denote the set of websites carrying the website-majority signal, then we define the degree of belief polarization of \mathcal{E} as:

$$\mathcal{BP}(\mathcal{E}) = |\hat{\rho}_{N,K}^A - \hat{\rho}_{N,K}^B|.$$

We say environment \mathcal{E} exhibits stronger belief polarization than \mathcal{E}' , if $\mathcal{BP}(\mathcal{E}) > \mathcal{BP}(\mathcal{E}')$.

When individuals are separated according to the (ex ante) accuracy of their signals, a larger personalization parameter λ will lead to more polarization as measured by \mathcal{BP} .²⁴

²³On Google’s search algorithm prioritizing websites reporting false information, see: “Harsh truths about fake news for Facebook, Google and Twitter”, *Financial Times*, November 21, 2016; “Google, democracy and the truth about internet search”, *The Observer*, December 4, 2016. On Facebook showing websites reporting false information in the top list of its trending topics, see: “Facebook fact-check: from Hillary’s health to 9/11, the latest lies we read,” *The Guardian*, September 9, 2016; “Three days after removing human editors, Facebook is already trending fake news,” *The Washington Post*, August 29, 2016.

²⁴Notice that, if individuals in different groups were to face also different initial rankings, ($r_1^A \neq r_1^B$), then this different rankings would clearly contribute to further accentuating the evolution of rankings seen by the two groups.

Proposition 3. *Let \mathcal{E}_λ be a personalized search environment with $\alpha > 0$, $\gamma > 0$, and with personalization parameter λ . Suppose there are two groups of individuals A and B of equal size and that have different levels of accuracy of their private signals, $p_A \neq p_B$. Then $\mathcal{BP}(\mathcal{E}_\lambda)$ is increasing in λ .*

Therefore, non-trivial personalization ($\lambda > 0$) can lead to different information held by individuals and hence to polarization of opinions. This result is in line with Flaxman *et al.* (2013), who show that search engines can lead to a relatively high level of ideological segregation, due to web search personalization embedded in the search engine’s algorithm and to individuals’ preference for like-minded sources of news. It is also in line with existing claims and empirical evidence suggesting that the Internet—together with the online platforms embedded in it—generally contributes towards increasing ideological segregation (Putnam, 2001; Sunstein, 2009; Pariser, 2011; Halberstam and Knight, 2014; Bessi *et al.*, 2015; Bar-Gill and Gandal, 2017). Most importantly, consistent with empirical evidence on the overall pattern of website traffic (Hindman, 2009; Gentzkow and Shapiro, 2011), our results show that such ideological segregation can also coexist with patterns of website traffic characterized by *concentration at the top* (Corollary 1).

In Section 6 we address the less obvious question of how personalization affects efficiency. As we will see, when agents have common value preferences, as assumed so far, and one of the two groups is “better informed” than the other, then personalization may decrease *ex ante* efficiency.

6 Efficiency and Asymptotic Learning

To assess the efficiency implications of the opinion dynamics studied in previous sections, we consider the search environments from an *ex-ante* and from an *interim* perspective. In line with the literature on social herding and social learning, we evaluate efficiency in terms of asymptotic learning (Banerjee, 1992; Bikhchandani *et al.*, 1992, 1998; Smith and Sørensen, 2000; Acemoglu *et al.*, 2011; Acemoglu and Ozdaglar, 2011). As a measure of efficiency, we use the limit probability of individual n (as $n \rightarrow \infty$) choosing the action corresponding to the true state of the world ($a_n = \omega$). At the interim stage, given the realization of true state of the world and of the signals of the websites, the probability of individual n clicking on one of the L websites reporting a correct signal can be written as:

$$\rho_{n,L} = \begin{cases} 0 & \text{if } L = 0 \\ \rho_{n,M \setminus K} & \text{if } 0 < L < \frac{M}{2} \\ \rho_{n,K} & \text{if } \frac{M}{2} \leq L < M \\ 1 & \text{if } L = M \end{cases} \quad (10)$$

where we recall that $K \subset M$ denotes the set of websites reporting the website-majority signal, so that $\rho_{n,K}$ ($\rho_{n,M \setminus K}$) denotes the probability of individual n reading a website reporting a website-majority (website-minority) signal, with $\rho_{n,M \setminus K} = 1 - \rho_{n,K}$. Accordingly, we can define a measure

of *interim efficiency* (\mathcal{P}_L), conditional on the number of websites reporting the correct signal (L):

$$\mathcal{P}_L(\alpha, \gamma; p) = \rho_{\infty, L} = \lim_{n \rightarrow \infty} \rho_{n, L}. \quad (11)$$

This allows us to define our measure of *ex ante efficiency* (\mathcal{P}) as follows:

$$\mathcal{P}(\alpha, \gamma; p, q) = \sum_{L=0}^M \binom{M}{L} q^L (1-q)^{M-L} \mathcal{P}_L(\alpha, \gamma; p). \quad (12)$$

As explained in Appendix A, we compute the limit ranking and clicking probabilities using the mean dynamics approximation (Norman, 1972; Izquierdo and Izquierdo, 2013). Essentially this involves fixing an interim search environment (essentially characterized by a number of websites L carrying the correct signal) and estimating the random clicking probabilities $\rho_{n, m}$ in Equation (3) for $m \in M$ by their expectations $\hat{\rho}_{n, m} = \mathbb{E}[\rho_{n, m}]$. This leads to deterministic recursions that are easily computed in the limit by means of ordinary differential equations. To obtain the ex ante efficiency, we take expectations over all interim environments as specified in Equation (12).

6.1 Comparative Statics

6.1.1 Interim Efficiency

In order to analyze ex-ante efficiency, we need to first study interim efficiency. The following proposition illustrates how interim efficiency changes as the number of websites carrying the correct signal L varies.

Proposition 4. *Let \mathcal{E} be a search environment with a uniform initial ranking r_1 and with N and κ large. Let $\langle \omega; (L, (y_m)) \rangle$ be an interim realization of \mathcal{E} , where L is the number of websites with the correct signal. Then, the interim efficiency is non-monotonic in L . Specifically:*

1. \mathcal{P}_L is always higher for $L \geq \frac{M}{2}$ than for $L < \frac{M}{2}$.
2. \mathcal{P}_L is weakly decreasing in L for $0 < L < \frac{M-1}{2}$ and for $\frac{M}{2} \leq L < M$.

Essentially, this proposition points out that: interim efficiency (\mathcal{P}_L) is greater when a majority of websites reports the correct signal than when a minority does so; the *advantage of the fewer* creates non-monotonicity in interim efficiency. Indeed, while it is always better to have a majority of websites reporting the correct signal, fewer websites carrying the correct signal can have an overall greater expected probability of being visited than if there are more (as suggested by Proposition 3). A graphical illustration of this result is shown in Figure 1, which plots interim efficiency as a function of L (for $p = 0.55$, a uniform initial ranking and for different values of α and γ).²⁵

²⁵When $L = 0$ (or $L = M$), all agents always access a website with a wrong (correct) signal. When $0 < L < \frac{M}{2}$, a minority of the websites carry the correct signal. Hence, for any given L in this range, an increase in γ monotonically increase interim efficiency. Instead, when $L \geq \frac{M}{2}$, a majority of the websites carry the correct signal. Hence, for any given L in this range, an increase in γ monotonically decreases interim efficiency.




Figure 1: *Interim* efficiency (\mathcal{P}_L) as a function of the number of websites with correct signal (L) for $\gamma = 0.5$ (black line) and other values of γ (dashed lines) and for two values of the search cost, $\alpha = 0$ (left) and $\alpha = 1$ (right). In both panels, $p = 0.55$, $M = 20$, and r_1 is uniform.

6.1.2 Ex Ante Efficiency

The following proposition shows that while a higher level of accuracy of the individuals' signals (p) and a lower degree of preference for like-minded news (γ) generally improve ex ante efficiency, this is not necessarily the case with the accuracy of the websites' signals (q) and the search cost (α). Depending on the degree of preference for like-minded news, q and α may have a positive or negative effect on ex ante efficiency.

Proposition 5. *Let \mathcal{E} be a search environment with a uniform initial ranking r_1 and with N and κ large. Then, ex ante efficiency is weakly increasing in p and weakly decreasing in γ . Moreover, there exists $0 < \bar{\gamma} < \frac{1}{2}$ such that ex ante efficiency is weakly increasing in q for $\gamma \in [0, \bar{\gamma}] \cup [1 - \bar{\gamma}, 1]$ and is weakly decreasing in α for $\gamma \in [1 - \bar{\gamma}, 1]$. In the remaining cases, ex ante efficiency can be decreasing in q for some $\gamma \in (\bar{\gamma}, 1 - \bar{\gamma})$ and can be increasing in α for some $\gamma \in [0, 1 - \bar{\gamma})$.*

A higher p always increases ex ante efficiency. This is not just the consequence of the direct (and trivial) effect of a higher probability of individual n having a correct private signal. The static effect is compounded by the dynamic one due to fact that the more individuals receive a correct signal, the higher the ranking of websites reporting a correct signal and, in turn, the higher the probability that other (subsequent) individuals will also choose a website reporting a correct signal. Indeed, even though individuals are assumed to be naïve with respect to the algorithm, the popularity component of the ranking algorithm combined with the search cost and preference for like-minded news create a mechanism whereby the private signals received by the individuals are implicitly aggregated and embedded in the website rankings. Corollary 2 makes this even clearer.

A higher γ decreases the probability of choosing a website reporting a correct signal when $L \geq \frac{M}{2}$, and increases it when $L < \frac{M}{2}$. Hence, in our framework with $q > p > \frac{1}{2}$ and naïve agents who do not learn from the ranking per se, a higher γ necessarily leads to a lower probability of asymptotic learning.

On the other hand, somewhat remarkably, a higher level of accuracy of websites' signals need not increase ex ante efficiency. This is true for intermediate levels of γ and is related to the *advantage*




Figure 2: *Ex ante* efficiency (\mathcal{P}) as a function of the preference for like-minded news (γ) for $q = 0.75$ for two values of the search cost, $\alpha = 1$ (black line), $\alpha = 0$ (dashed line) (left); and for $\alpha = 1$ and two values of website accuracy, $q = 0.75$ (black line), $q = 0.85$ (dashed line) (right). In both panels, $p = 0.55$, $M = 20$, and r_1 is uniform.

of the fewer property discussed in Section 4. Higher values of q make it more likely that the number of websites with incorrect minority signal is small and the ones with correct signal is large. Because of the *advantage of the fewer* effect, this increases the clicking probability on such websites while reducing it for websites with the correct majority signal, thus reducing ex ante efficiency.

Finally, a higher search cost α can increase or decrease ex ante efficiency depending on the value of γ . When γ is small, it tends to enhance the ranking of websites carrying the website-majority signal, which tends to increase ex ante efficiency. When γ is large, a higher α tends to enhance the ranking of websites carrying the website-minority signal, which tends to reduce the ex ante probability of asymptotic learning. A graphical illustration of this is shown in Figure 2 which plots ex ante efficiency as a function of γ for $\alpha = 0$ (equivalent to randomized ranking) and for $\alpha = 1$ (left panel). It also shows ex ante efficiency as a function γ for $q = 0.75$ and $q = 0.85$.

It is insightful to compare the case of rankings that are endogenous to individuals' searches ($\kappa < \infty$ or $\nu < 1$) to ones where the ranking is random and uniform throughout (r_1 uniform and $\kappa = \infty$). It is not difficult to see that in this case ex ante efficiency is as in a search environment where search cost is zero, so that individuals' website choices correspond to *ranking-free* ones. The following result then follows from the proof of Proposition 5.

Corollary 2. *Let \mathcal{E} be a search environment with a uniform initial ranking r_1 , and let \mathcal{E}' be a search environment where the initial ranking and every subsequent ranking is uniform ($\kappa = \infty$). Let \mathcal{P} and \mathcal{P}' denote the ex ante efficiency of \mathcal{E} and \mathcal{E}' respectively. Then there exists $0 < \bar{\gamma} \leq 1$ such that $\mathcal{P} \geq \mathcal{P}'$ for $\gamma \in [0, \bar{\gamma}]$, while $\mathcal{P} \leq \mathcal{P}'$ for $\gamma \in [\bar{\gamma}, 1]$.*

This shows that a fixed random ranking is inferior in terms of ex ante efficiency to a popularity driven one, unless there is a sufficiently strong preference for like-minded news, in which case a random ranking is efficiency-enhancing.

6.2 Ex Ante Efficiency and Personalization

This section discusses the case of personalization of search results, where individuals in different groups have different levels of accuracy (say p_A and p_B). Personalization here can be seen as progressively “separating” the two groups thereby progressively uncoupling their rankings and so switching off potential externalities from one group to the other. To the extent that the high accuracy group exerts a positive externality on the groups’ rankings and overall ex ante efficiency, increasing the personalization may inhibit overall ex ante efficiency.

Proposition 6. *Let \mathcal{E}_λ be a personalized search environment with personalization parameter λ , a uniform initial ranking r_1 and with N and κ large. Suppose there are two groups of individuals A and B of equal size that have different levels of accuracy of their private signals, say $\frac{1}{2} < p_A < p_B < \bar{p}$, for some $\bar{p} < 1$. Then, there exists $\frac{1}{2} < \bar{p} < 1$ such that ex-ante efficiency is weakly decreasing in λ .*

The proposition shows how personalization can contribute towards lower ex ante efficiency. It suggests that individuals with higher accuracy provide an overall positive externality on the probability of clicking on a website with the correct signal when they are mixed with individuals of lower accuracy, as is the case when there is less personalization.

Formally, the negative effect of personalization on ex ante efficiency shows up through the concavity of the ex ante efficiency function (\mathcal{P}^λ) with respect to the accuracy parameter (p) for any λ . The ex ante efficiency of a non-personalized ranking ($\lambda = 0$) corresponds to the ex ante efficiency of a single group with a level of accuracy that is an average of p_A and p_B ; whereas the ex ante efficiency of a personalized ranking ($\lambda > 0$) corresponds to an average of the ex ante efficiency of two groups, each with its own level of accuracy. Concavity of the ex ante efficiency measure with respect to p means that the latter ($\mathcal{P}^\lambda, \lambda > 0$) is below the former ($\mathcal{P}^\lambda, \lambda = 0$), and this carries over to any comparison for any given pair of λ 's. The intuition for why \mathcal{P}^λ is concave in p has to do with the fact that increasing p increases interim efficiency \mathcal{P}_L^λ , for any L , but where the marginal effect is increasing in p when L represents a minority ($L = M \setminus K$) and is decreasing in p when it represents a majority ($L = K$) of websites. Because the latter cases where the websites with a correct signal represent a majority of websites are more likely, since $q > \frac{1}{2}$, when the accuracy levels of the individuals are not too large, the negative effects can be shown to dominate the positive ones. This makes the overall function concave.

Finally, we should emphasize, that our analysis focuses on common value preferences. As we will see in Section 7.2, when preferences vary across individuals, then there can be more scope for personalized ranking algorithms to be efficiency-enhancing.

7 Extensions

We briefly explore two extensions of the model.

7.1 Domain Bias – M'

We now revisit the assumption made in Section 2.3 that agents form their ranking-free website choices from observing the headlines of $M' \subset M$ websites, where it was assumed that $M' = M$. We now relax that assumption to show that the main properties derived above not only continue to go through in the more realistic case, where M' is a strict subset of M , but in some cases they can lead to significantly more severe consequences. In particular, assuming that agents form their ranking-free website choices from observing the headlines of $M' \neq M$ websites introduces an additional behavioral bias. Such a behavioral bias would be consistent with what Tversky and Kahneman (1974) refer to as an availability bias due to the retrievability of instances. For the sake of simplifying the exposition, we simply refer to such a behavioral preference as a *domain bias*.

Assume, for simplicity, that there is no personalization ($\lambda = 0$), no preference for like-minded news ($\gamma = 0$), and that M' consists of the M' websites with the highest ranking probability, that is, websites such that $r_{n,m} \leq r_{n,m'}$ for any $m \in M \setminus M', m' \in M'$ (e.g., only websites in the first page of search results). Given M' , let $y_{K'} \in \{0, 1\}$ denote the signal that is reported by a majority of websites in M' , the M' website-majority signal (i.e., the most frequent signal reported in the websites' *snippets* provided in the first page of search results); let also $K' \leq M'$ denote the number of websites that report that signal, that is, $K' = \#\{m \in M' \mid y_m = y_{K'}\}$. Accordingly, we define as $J \leq N$ the number of individuals receiving a private signal equivalent to their website-majority signal, that is $J = \#\{n \in N \mid x_n = y_{K'}\}$.

Since $q > p$, agents' preference vector ρ_n^* now takes the form:

$$\rho_{n,m}^* = \frac{\mathbb{I}_{\{y_m = y_{K'}, m \in M'\}}}{K'} = \begin{cases} \frac{1}{K'} & \text{if } y_m = y_{K'} \text{ and } m \in M' \\ 0 & \text{else .} \end{cases} \quad (13)$$

In particular, agents only consult websites in M' . As a result, all websites in $M \setminus M'$ will have a choice probability of zero, $\rho_{n,m} = 0$ for any $m \in M \setminus M'$ and such websites will tend to also receive zero ranking probability for N large. While all the results derived above go through in this case, after appropriate adaptations, the main difference lies in the fact that what drives the opinion dynamics and the choices of individuals is not the majority signal y_M taken from all websites, but a majority signal $y_{M'}$ taken from just a subsample. This means that its informational content may be significantly weaker depending on the relative size of M' . Moreover, for $\lambda > 0$, differences in M' across individuals may lead to a different perceived website-majority signals. In turn, this may affect the individual choice over websites and, ultimately, reinforce *belief polarization*.^{26,27}

²⁶It is worth noticing that the search cost combined with a *domain bias*, may provide a rationale to explain the large size of the SEME effect found in Epstein and Robertson (2015). Future empirical research may attempt to disentangle the impact of these two components (i.e., search cost and *domain bias*) on individual choice over websites.

²⁷If one studies ranking algorithms from a mechanism design perspective, then a smaller $M' \subsetneq M$ can be advantageous for the principal (i.e., the designer of the ranking algorithm), as it gives more possibilities to steer traffic towards a narrower set of websites. Notice also that the ranking algorithm can potentially learn a lot about the

7.2 Heterogeneous Preferences and Personalization

Our basic framework focuses on the case where individuals have homogeneous preferences over a common state of the world (e.g., side effects of a vaccine). In Section 6, we saw how search personalization may be detrimental for efficiency in such a context. As the following example points out, when individuals have heterogeneous preferences (e.g., they care differently about different attributes of a good or product), and this is reflected in the personalization of search results, then a personalized ranking algorithm may be efficiency enhancing relative to a non-personalized one.

Example 1. (Personalized ranking with private values). Consider two groups A and B , and a two-dimensional state variable $\omega = (\omega^A, \omega^B) \in \{0, 1\}^2$, where ω^A, ω^B are i.i.d. $(\frac{1}{2}, \frac{1}{2})$ Bernoulli random variables, which take one of two values in $\{0, 1\}$. Individuals in different groups differ with respect to their utility function:

$$u_n^\ell(a_n, \omega) = \begin{cases} 1 & \text{if } a_n = \omega^\ell \\ 0 & \text{else,} \end{cases}$$

for $\ell = A, B$. As before, individuals and websites receive signals, however, now individuals in different groups differ with respect to the correlation of their private signal with the state variable ω . In particular, an agent in group ℓ receives a signal $x_n^\ell \in \{0, 1\}$, satisfying, for $\ell = A, B$:

$$\mathbb{P}(x_n^\ell = \omega^\ell | \omega^\ell = z) = p^\ell > \frac{1}{2}, \text{ for any } z \in \{0, 1\},$$

where we again assume $p^A = p^B = p$. In other words, we consider a pure private value model, where individuals in each group care only about the state variable concerning their group and receive a private signal only depending on such a state variable. This can be interpreted as a setting, where the two groups care only about a specific characteristic of the product or service and, accordingly, they only receive a private signal regarding the characteristic that they care about (e.g., ω^A, ω^B may represent the atmosphere and the food quality of a restaurant, with individuals in group A just caring about atmosphere, while the ones in group B just care about the food quality).

On the other hand, each website m receives (and reports) a coarse signal regarding the two-dimensional state variable ω (e.g., the average quality – bad/good – of a restaurant which may be a combination of its atmosphere and of the food taste). In particular, suppose:

$$\mathbb{P}(y_m = 0 | \omega^A + \omega^B < 1) = q_0 \text{ and } \mathbb{P}(y_m = 1 | \omega^A + \omega^B \geq 1) = q_1,$$

where, as before, we assume $q_0, q_1 > p > \frac{1}{2}$, are sufficiently large such that the coarse signals of a website are always informative for individuals in both groups A and B . Individuals in both groups will therefore rationally revise upwards their beliefs on ω^A or ω^B being equal to one whenever they

environment. Kremer *et al.* (2014) take such an informed principle approach in the context of recommender systems. We leave the study of this within a context of ranking algorithms for future work.

read a website reporting signal $y_m = 1$, and will revise them downwards whenever they read a website reporting signal $y_m = 0$.²⁸

Hence, as in the case where all individuals share the same preferences over the state of the world, an algorithm personalizing search results may lead to polarized beliefs. However, unlike the case of common-value preferences, since neither group has a positive information externality on the other, personalization is more efficient when the two groups have uncorrelated preferences. \square

Together with Proposition 6, this example shows how personalization is likely to be efficient in some contexts (e.g., with private values) and inefficient in others (e.g., with common values). This is consistent with the intuition expressed in Lazer (2015), p. 1090, that “Social algorithms are often quite helpful; when searching for pizza in Peoria, it helps not to get results about Famous Ray’s in Manhattan. However, personalization might not be so benign in other contexts, raising questions about equity, justice, and democracy.” It is important for a ranking algorithm to be able to decide on the “right” degree of personalization, based on just the entries of the query. This may not always be a straightforward to accomplish (Hannak *et al.*, 2013).

8 Conclusions

The economics literature has largely overlooked the informational role played by algorithms used, for example, by search engines and social media to rank the information provided to their users. To address this neglected issue, we have developed an empirically-grounded model that allows us to study how basic features of a search engine’s algorithm interact with individual search behavior and how this affects website traffic, opinion dynamics and asymptotic learning.

The results inform both the popular and academic debate regarding the effect of the Internet’s architecture on market concentration and ideological segregation. While some scholars argue that the Internet has led to an increased concentration of the media sector (McChesney, 2013), others suggest that the Internet may actually have “too little” concentration from favoring self-segregation and “cyber-balkanization” (Putnam, 2001; Sunstein, 2009; Halberstam and Knight, 2014; Bessi *et al.*, 2015). Consistent with empirical evidence on the pattern of website traffic (Hindman, 2009; Gentzkow and Shapiro, 2011) and on the role of algorithmic gatekeepers (Flaxman *et al.*, 2013), our analysis shows that these two arguments are not necessarily mutually exclusive. Patterns of websites traffic characterized by *concentration at the top* may well co-exist with *belief polarization* in opinion dynamics.

Most importantly, the main insights of the paper point out that—compared to a situation where individuals observe randomly ranked websites—popularity driven rankings have an overall

²⁸Notice that $\mathbb{P}(\omega^A = 0 | y_m = 0) = \mathbb{P}(\omega^B = 0 | y_m = 0) = \frac{q_0 + (1 - q_1)}{q_0 + 3(1 - q_1)}$, which, for $q_0, q_1 > p$, is higher than $\mathbb{P}(\omega^A = 0 | x_n^A = 0) = \mathbb{P}(\omega^B = 0 | x_n^B = 0) = p$. At the same time, $\mathbb{P}(\omega^A = 1 | y_m = 1) = \mathbb{P}(\omega^B = 1 | y_m = 1) = \frac{2q_1}{3q_1 + (1 - q_0)}$ which, for $q_0, q_1 > p$, is higher than $\mathbb{P}(\omega^A = 1 | x_n^A = 1) = \mathbb{P}(\omega^B = 1 | x_n^B = 1) = p$.

positive effect on the efficiency of individual choices, and enhance the aggregation of the private information dispersed among individuals and websites. At the same time, ranking algorithms can also contribute to the diffusion of misinformation. The lower the ex-ante quality of information of minority websites, the higher the overall share of traffic that they attract (*advantage of the fewer*). The model further provides insights on a controversial component of the ranking algorithm: the personalization of search results. While it might be useful for search queries on private value issues (e.g., attributes of a commercial product), personalization hinders asymptotic learning when individuals are looking for information on common value issues (e.g., side effects of a vaccine).

So far, public authorities have worried about the possible strategic manipulation of website rankings by search engines regarding transactional queries.²⁹ Our results suggest that, even in the “best scenario,” where search engines do not have any incentive to strategically manipulate their search results to increase their advertising profits (as in De Cornière and Taylor 2014; Taylor 2013; Burguet *et al.* 2015; Hagiu and Jullien 2014) or to influence election outcomes (as in Epstein and Robertson 2015), there might still be subtler distortions in the pattern of website traffic and opinion dynamics arising from the interaction of the ranking algorithm with individuals’ online search behavior. Accordingly, our theoretical framework provides policy regulators with both a rationale to request information on the architecture of the ranking algorithms as well as some guidance on how to evaluate them in terms of website traffic and efficiency.³⁰ In general, identifying the key components of ranking algorithms, enhancing their transparency and promoting individuals’ awareness of such components and of the dynamics at play, might prove useful.³¹

²⁹http://europa.eu/rapid/press-release_MEMO-15-4781_en.htm

³⁰For example, the potential presence of a *rich-get-richer* dynamic highlighted by our model may be of interest to competition authorities since it can lead to large asymmetries in market size (share of traffic) of otherwise identically informative websites. As for personalization and *belief polarization*, a “dirigiste” approach might suggest a further role for regulation authorities in supervising the architecture of ranking algorithms to reduce the extent of personalization in search queries involving common-value issues (e.g., side effects of vaccines). A less paternalistic approach might simply consist in encouraging search engines to allow each user to choose between search results based on previous search queries of individuals with “similar characteristics” and non-personalized results. In this way, each user might endogenously choose whether or not to personalize her search results for given queries.

³¹Survey evidence on a subsample of Facebook’s users reveal that above the 60% of respondents are not aware of even the existence of the News Feed algorithm (Eslami *et al.*, 2015).

References

- ACEMOGLU, D., DAHLEH, M. A., LOBEL, I. and OZDAGLAR, A. (2011). Bayesian Learning in Social Networks. *The Review of Economic Studies*, **78**, 1201–1236.
- and OZDAGLAR, A. (2011). Opinion Dynamics and Learning in Social Networks. *Dynamic Games and Applications*, **1** (1), 3–49.
- ALAOU, L. and GERMANO, F. (2016). Time Scarcity and the Market for News. *UPF Working Paper*.
- ALLCOTT, H. and GENTZKOW, M. (2017). *Social Media and Fake News in the 2016 Election*. Working Paper 23089, National Bureau of Economic Research.
- ALTMAN, A. and TENNENHOLTZ, M. (2008). Axiomatic foundations for ranking systems. *Journal of Artificial Intelligence Research*, **31**, 473–495.
- BAKSHY, E., MESSING, S. and ADAMIC, L. (2015). Exposure to Ideologically Diverse News and Opinion on Facebook. *Science*.
- BANERJEE, A. (1992). A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, **3** (107), 797–817.
- BAR-GILL, S. and GANDAL, N. (2017). *Online Exploration, Content Choice & Echo Chambers: An Experiment*. Tech. rep., CEPR Discussion Papers.
- BESSI, A., COLETTI, M., DAVIDESCU, G. A., SCALA, A. and CALDARELLI, G. (2015). Science vs Conspiracy: Collective Narratives in the Age of Misinformation. *PloS one*, **10** (2), 2.
- BIKHCHANDANI, S., HIRSHLEIFER, D. and WELCH, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, **100**, 992–1026.
- , — and — (1998). Learning From the Behavior of Others: Conformity, Fads, and Informational Cascades. *The Journal of Economic Perspectives*, pp. 151–170.
- BOHREN, J. A. (2016). Informational with Model Misspecification. *Journal of Economic Theory*, **163**, 222–247.
- BRIN, S. and PAGE, L. (1998). The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks and ISDN Systems*, **30** (1-7), 107–117.
- BURGUET, R., CAMINAL, R. and ELLMAN, M. (2015). In Google we Trust? *International Journal of Industrial Organization*, **39**, 44–55.
- CARVALHO, C., KLAGGE, N. and MOENCH, E. (2011). The Persistent Effects of a False News Shock. *Journal of Empirical Finance*, **18** (4), 597–615.
- CHAN, J. and SUEN, C. (2008). A Spatial Theory of News Consumption and Electoral Competition. *The Review of Economic Studies*, **75** (3), 699–728.
- CHITIKA (2013). *The Value of Google Result Positioning*.
- CHO, J. and ROY, S. (2004). Impact of search engines on page popularity. In *Proceedings of the 13th international conference on World Wide Web*, ACM, pp. 20–29.

- , — and ADAMS, R. E. (2005). Page Quality: In Search of an Unbiased Web Ranking. *SIGMOD*, **14**.
- DE CORNIÈRE, A. and TAYLOR, G. (2014). Integration and Search Engine Bias. *RAND Journal of Economics*, **45** (3), 576–597.
- DEAN, B. (2013). Google’s 200 Ranking Factors. *Search Engine Journal*, **May 31** (<http://www.searchenginejournal.com/infographic-googles-200-ranking-factors>).
- DELLAVIGNA, S. and GENTZKOW, M. (2010). Persuasion: empirical evidence. *Annu. Rev. Econ.*, **2** (1), 643–669.
- DEMANGE, G. (2012). On the Influence of a Ranking System. *Social Choice and Welfare*, **39** (2), 431–455.
- (2014). Collective Attention and Ranking Methods. *Journal of Dynamics and Games*, **1** (1), 17–43.
- DEMARZO, P. M., VAYANOS, D. and ZWIEBEL, J. (2003). Persuasion Bias, Social Influence, and Uni-Dimensional Opinions. *Quarterly Journal of Economics*, **118** (3), 909–968.
- EPSTEIN, R. (2013). Google’s Gotcha. 15 Ways Google Monitors You. *U.S. News & World Report*, (May 10, 2013).
- and ROBERTSON, R. E. (2015). The Search Engine Manipulation Effect (SEME) and its Possible Impact on the Outcomes of Elections. *Proceedings of the National Academy of Sciences*, **112** (33), E4512—E4521.
- ESLAMI, M., RICKMAN, A., VACCARO, K., ALEYASEN, A., VUONG, A., KARAHALIOS, K., HAMILTON, K. and SANDVIG, C. (2015). I Always Assumed that I Wasn’t Really that Close to [Her]: Reasoning about Invisible Algorithms in News Feeds. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, (153–162).
- EYSTER, E. and RABIN, M. (2010). Naïve Herding in Rich-Information Settings. *American Journal of Economics: Microeconomics*, **2** (4), 221–243.
- FLAXMAN, S., GOEL, S. and RAO, J. M. (2013). Ideological Segregation and the Effects of Social Media on News Consumption. *Working Paper*, **10**.
- FORTUNATO, S., FLAMMINI, A., MENCZER, F. and VESPIGNANI, A. (2006). Topical interests and the mitigation of search engine bias. *Proceedings of the National Academy of Sciences*, **103** (34), 12684–12689.
- GENTZKOW, M. and SHAPIRO, J. (2006). Media Bias and Reputation. *Journal of Political Economy*, **114** (2), 280–316.
- and SHAPIRO, J. M. (2010). What drives media slant? evidence from u.s. daily newspapers. *Econometrica*, **78** (1), 35–71.
- and — (2011). Ideological Segregation Online and Offline. *The Quarterly Journal of Economics*, **126** (4), 1799–1839.

- GLICK, M., RICHARDS, G., SAPOZHNIKOV, M. and SEABRIGHT, P. (2011). How Does Page Rank Affect User Choice in Online Search? *Working Paper*.
- GOLDMAN, E. (2006). Search Engine Bias and the Demise of Search Engine Utopianism. *Yale Journal of Law & Technology*, pp. 6–8.
- GRANKA, L. A. (2010). The Politics of Search: A Decade Retrospective. *The Information Society*, **26**, 364–374.
- GRIMMELMANN, J. (2009). The Google Dilemma. *New York Law School Law Review*, **53**, 939–950.
- GUAN, Z. and CUTRELL, E. (2007). An Eye Tracking Study of the Effect of Target Rank on Web Search. *CHI 2007 Proceedings Gaze & Eye Tracking*.
- GUARINO, A. and JEHIEL, P. (2013). Social Learning with Coarse Inference. *American Journal of Economics: Microeconomics*, **5** (1), 147–174.
- HAGIU, A. and JULLIEN, B. (2014). Search Diversion and Platform Competition. *International Journal of Industrial Organization*, **33**, 46–80.
- HALBERSTAM, Y. and KNIGHT, B. (2014). Homophily, Group Size, and the Diffusion of Political Information in Social Networks: Evidence from Twitter.
- HANNAK, A., SAPIEŻYŃSKI, P., KAKHKI, A. M., KRISHNAMURTHY, B., LAZER, D., MISLOVE, A. and WILSON, C. (2013). Measuring Personalization of Web Search. *Proceedings of the Twenty-Second International World Wide Web Conference (WWW'13)*.
- HARGITTAI, E. (2004). The Changing Online Landscape. *Community practice in the network society: local action/global interaction*.
- HAZAN, J. G. (2013). Stop Being Evil: A Proposal for Unbiased Google Search. *Michigan Law Review*, **111** (789), 789–820.
- HINDMAN, M. (2009). *The Myth of Digital Democracy*. Princeton University Press.
- HORRIGAN, J. B. (2006). The Internet as a Resource for News and Information about Science. *Pew Internet & American Life Project*.
- INTRONA, L. D. and NISSENBAUM, H. (2000). Shaping the web: Why the politics of search engines matters. *The information society*, **16** (3), 169–185.
- IZQUIERDO, S. S. and IZQUIERDO, L. R. (2013). Stochastic Approximation to Understand Simple Simulation Models. *Journal of Statistical Physics*, **151**, 254–276.
- JANSEN, B. J., BOOTH, D. L. and SPINK, A. (2008). Determining the Informational, Navigational, and Transactional Intent of Web Queries. *Information Processing and Management*, (44), 1251–1266.
- KATA, A. (2012). Anti-Vaccine Activists, Web 2.0, and the Postmodern Paradigm—An Overview of Tactics and Tropes used Online by the Anti-Vaccination Movement. *Vaccine*, **30** (25), 3778–3789.
- KEARNEY, M. S. and LEVINE, P. B. (2014). *Media Influences on Social Outcomes: The Impact of MTV's 16 and Pregnant on Teen Childbearing*. Working Paper 19795, National Bureau of Economic Research.

- KLIMAN-SILVER, C., HANNAK, A., LAZER, D., WILSON, C. and MISLOVE, A. (2015). Location, Location, Location: The Impact of Geolocation on Web Search Personalization. In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*, ACM, pp. 121–127.
- KREMER, I., MANSOUR, Y. and PERRY, M. (2014). Implementing the wisdom of the crowd. *Journal of Political Economy*, **122** (5), 988–1012.
- LAZER, D. (2015). The Rise of the Social Algorithm. *Science*, **348** (6239), 1090–1091.
- MCCHESNEY, R. W. (2013). *Digital Disconnect: How Capitalism is Turning the Internet Against Democracy*. The New Press.
- MENCZER, F., FORTUNATO, S., FLAMMINI, A. and VESPIGNANI, A. (2006). Googlearchy or Geoleocracy. *IEEE Spectrum Online*.
- MOCANU, D., ROSSI, L., ZHANG, Q., KARSAI, M. and QUATTROCIOCCHI, W. (2015). Collective Attention in the Age of (Mis) Information. *Computers in Human Behavior*, **51**, 1198–1204.
- MONZÓN, I. and RAPP, M. (2014). Observational Learning with Position Uncertainty. *Journal of Economic Theory*, **154**, 375–402.
- MOZ (2013). *2013 Search Engine Ranking Factors*. (<http://moz.com/search-ranking-factors>).
- MUELLER-FRANK, M. (2013). A General Framework for Rational Learning in Social Networks. *Theoretical Economics*, **8**, 1–40.
- MULLAINATHAN, S. and SHLEIFER, A. (2005). The Market for News. *American Economic Review*, **95** (4), 1031–1053.
- NAPOLI, P. M. (2015). Social media and the public interest: Governance of news platforms in the realm of individual and algorithmic gatekeepers. *Telecommunications Policy*, **39** (9), 751–760.
- NORMAN, M. F. (1972). *Markov Processes and Learning Models*. Academic Press.
- NOVARESE, M. and WILSON, C. (2013). Being in the Right Place: A Natural Field Experiment on List Position and Consumer Choice. *Working Paper*.
- PALACIOS-HUERTA, I. and VOLIJ, O. (2004). The measurement of intellectual influence. *Econometrica*, **72** (3), 963–977.
- PAN, B., HEMBROOKE, H., JOACHIMS, T., LORIGO, L., GAY, G. and GRANKA, L. (2007). In Google We Trust: Users’ Decisions on Rank, Position, and Relevance. *Journal of Computer-Mediated Communication*, **12**, 801–823.
- PARISER, E. (2011). *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. Penguin Books.
- PRAT, A. and STRÖMBERG, D. (2013). The Political Economy of Mass Media. In: *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress*.
- PUTNAM, R. D. (2001). *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster.

- RABIN, M. and SCHRAG, J. L. (1999). First Impressions Matter: A Model of Confirmatory Bias. *The Quarterly Journal of Economics*, pp. 37–82.
- RIEDER, B. (2005). Networked control: Search engines and the symmetry of confidence. *International Review of Information Ethics*, **3** (1), 26–32.
- and SIRE, G. (2013). Conflict of Interest and the Incentives to Bias: A Microeconomic Critique of Google’s Tangled Position on the Web. *New Media & Society*, **0**, 1–17.
- SHAO, C., CIAMPAGLIA, G. L., FLAMMINI, A. and MENCZER, F. (2016). Hoaxy: A Platform for Tracking Online Misinformation. In *Proc. Third Workshop on Social News On the Web (WWW SNOW)*.
- SMITH, L. and SØRENSEN, P. (2000). Pathological Outcomes of Observational Learning. *Econometrica*, **68** (2), 371–398.
- SOBBRIO, F. (2014). Citizen-Editors’ Endogenous Information Acquisition and News Accuracy. *Journal of Public Economics*, **113**.
- SPINK, A. and JANSEN, B. J. (2004). *Web Search: Public Searching of the Web*. Kluwer Academic Publishers.
- STRÖMBERG, D. (2004). Mass Media Competition, Political Competition, and Public Policy. *Review of Economic Studies*, **11** (1).
- SUNSTEIN, C. R. (2009). *Republic.com 2.0*. Princeton University Press.
- TAYLOR, G. (2013). Search Quality and Revenue Cannibalization by Competing Search Engines. *Journal of Economics & Management Strategy*, **22** (3), 445–467.
- TUFEKCI, Z. (2015). Algorithmic Harms beyond Facebook and Google: Emergent Challenges of Computational Agency. *J. on Telecomm. & High Tech. L.*, **13**, 203.
- TVERSKY, A. and KAHNEMAN, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *science*, **185** (4157), 1124–1131.
- VAUGHN, A. (2014). *Google Ranking Factors. SEO Checklist*. <http://www.vaughns-1-pagers.com/internet/google-ranking-factors.htm>.
- WHITE, R. (2013). Beliefs and Biases in Web Search. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, ACM, pp. 3–12.
- WHITE, R. W. and HORVITZ, E. (2015). Belief dynamics and biases in web search. *ACM Transactions on Information Systems (TOIS)*, **33** (4), 18.
- XING, X., MENG, W., DOOZAN, D., FEAMSTER, N., LEE, W. and SNOEREN, A. C. (2014). Exposing Inconsistent Web Search Results with Bobble. In *Passive and Active Measurement*, Springer, pp. 131–140.
- YOM-TOV, E., DUMAIS, S. and GUO, Q. (2013). Promoting Civil Discourse Through Search Engine Diversity. *Social Science Computer Review*, pp. 1–10.

APPENDIX

A Mean Dynamics Approximation

In order to evaluate the actions taken by an agent in the limit as $N \rightarrow \infty$, we use some techniques of stochastic approximation from Norman (1972) as exposed in Izquierdo and Izquierdo (2013), which we refer to as the *mean dynamics approximation*. We here give a brief outline in order to follow our calculations and proofs, but we refer to the latter two sources for more details. The basic idea of the approach is to use the expected increments to evaluate the long run behavior of a dynamic process with stochastic increments. Rewrite the ranking probabilities as,

$$\begin{aligned}
 r_{t,m} &= \nu_t r_{t-1,m} + (1 - \nu_t) \rho_{t-1,m} \\
 &= \frac{\kappa_t}{1 + \kappa_t} r_{t-1,m} + \frac{1}{1 + \kappa_t} \rho_{t-1,m} \\
 &= r_{t-1,m} + \frac{1}{1 + \kappa_t} (\rho_{t-1,m} - r_{t-1,m}) \\
 &= r_{t-1,m} + \frac{1}{1 + \kappa_t} \left(\frac{(r_{t-1,m})^\alpha \cdot \rho_{t-1,m}^*}{\sum_{m'} (r_{t-1,m'})^\alpha \cdot \rho_{t-1,m'}^*} - r_{t-1,m} \right),
 \end{aligned}$$

where in order to obtain sharper convergence results, we let κ_t and hence $\nu_t \in (0, 1)$ vary with t . It is clear that the only stochastic term is in given by the expressions $\rho_{t-1,m}^*$. Replacing these with their expectations yields the deterministic recursion in $\hat{r}_{t,m}$

$$\hat{r}_{t,m} = \hat{r}_{t-1,m} + \frac{1}{1 + \kappa_t} (\hat{\rho}_{t-1,m} - \hat{r}_{t-1,m})$$

where

$$\hat{\rho}_{t-1,m} = \mathbb{E}[\rho_{t-1,m}] = p \frac{(\hat{r}_{t-1,m})^\alpha \cdot \hat{\rho}_m^0}{\sum_{m'} (\hat{r}_{t-1,m'})^\alpha \cdot \hat{\rho}_{m'}^0} + (1 - p) \frac{(\hat{r}_{t-1,m})^\alpha \cdot \hat{\rho}_m^1}{\sum_{m'} (\hat{r}_{t-1,m'})^\alpha \cdot \hat{\rho}_{m'}^1} \quad (14)$$

is the expected clicking probability of the individual entering in period $t - 1$, and where

$$\hat{\rho}_m^0 = \begin{cases} \frac{\gamma}{L} & \text{if } m \in L, 0 < L < \frac{M}{2} \\ \frac{1}{L} & \text{if } m \in L, \frac{M}{2} \leq L < M \\ \frac{1-\gamma}{M-L} & \text{if } m \notin L, 0 < L < \frac{M}{2} \\ 0 & \text{if } m \notin L, \frac{M}{2} \leq L < M \end{cases}, \quad \hat{\rho}_m^1 = \begin{cases} 0 & \text{if } m \in L, 0 < L < \frac{M}{2} \\ \frac{1-\gamma}{L} & \text{if } m \in L, \frac{M}{2} \leq L < M \\ \frac{1}{M-L} & \text{if } m \notin L, 0 < L < \frac{M}{2} \\ \frac{\gamma}{M-L} & \text{if } m \notin L, \frac{M}{2} \leq L < M. \end{cases}$$

Here $\hat{\rho}_m^0$ represents the expected probability of an individual choosing a website m , contingent on having received a correct signal, absent any search cost. Analogously, $\hat{\rho}_m^1$ represents the expected probability of an individual choosing a website m , contingent on having received an incorrect signal, absent any search cost.³² $\hat{\rho}_m^0$ and $\hat{\rho}_m^1$ are fixed coefficients that do not vary with t . In particular, in order to apply the basic approximation theorem we assume κ_t is of the order $O(t)$ so that, $\frac{1}{1+\kappa_t} \rightarrow 0$, and to guarantee smoothness and avoid boundary problems, we assume there exists $\epsilon > 0$ such that, each $\hat{r}_{t,m} \geq \epsilon$ for all t, m . Moreover, replacing $\hat{r}_{t-1,m}$ with x_m (and hence the vector \hat{r}_{t-1} with the vector $x = (x_1, \dots, x_M)$), this yields a function $g : \Delta_\epsilon(M) \rightarrow \mathbb{R}^M$, defined for

³²Clearly, $\hat{\rho}_m^0 = \hat{\rho}_m^1 = 0$ if $m \in L, L = 0$ or $m \notin L, L = M$. At the same time, $\hat{\rho}_m^0 = \hat{\rho}_m^1 = 1$ if $m \in L, L = M$ or $m \notin L, L = 0$.

$m = 1, 2, \dots, M$, by,

$$g_m(x) = \mathbb{E}[\rho_{t-1,m} - r_{t-1,m} \mid r_{t-1,m} = x_m \text{ for } m = 1, \dots, M] = \theta_m(x) - x_m,$$

where $\theta : \Delta_\epsilon(M) \rightarrow \mathbb{R}^M$ is defined by,

$$\theta_m(x) = p \frac{(x_m)^\alpha \cdot \widehat{\rho}_m^*{}^0}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{\rho}_{m'}^*{}^0} + (1-p) \frac{(x_m)^\alpha \cdot \widehat{\rho}_m^*{}^1}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{\rho}_{m'}^*{}^1}.$$

Given that the function g is smooth in x on $\Delta_\epsilon(M)$, it can be shown that the expected limit of our stochastic process can be obtained by solving the ordinary differential equation $\dot{x} = g(x)$. In particular, for any given initial condition r_0 , there is a unique limit, and for large enough values of κ the stochastic process r_t tends to follow the unique solution trajectory of the differential equation $\dot{x} = g(x)$.

B Proofs

Proof of Proposition 1. We prove directly the case with $\gamma \geq 0$. The expected ranking probabilities are then given by:

$$\begin{aligned} \widehat{r}_{n,m} &= \nu \widehat{r}_{n-1,m} + (1-\nu) \widehat{\rho}_{n-1,m} \\ &= \nu \widehat{r}_{n-1,m} + (1-\nu) \left(p \frac{(\widehat{r}_{n-1,m})^\alpha \cdot \widehat{\rho}_m^*{}^0}{\sum_{m'} (\widehat{r}_{n-1,m'})^\alpha \cdot \widehat{\rho}_{m'}^*{}^0} + (1-p) \frac{(\widehat{r}_{n-1,m})^\alpha \cdot \widehat{\rho}_m^*{}^1}{\sum_{m'} (\widehat{r}_{n-1,m'})^\alpha \cdot \widehat{\rho}_{m'}^*{}^1} \right) \\ &= \nu \widehat{r}_{n-1,m} + (1-\nu) \left(\frac{p \cdot \widehat{\rho}_m^*{}^0}{\sum_{m'} (\widehat{r}_{n-1,m'})^\alpha \cdot \widehat{\rho}_{m'}^*{}^0} + \frac{(1-p) \cdot \widehat{\rho}_m^*{}^1}{\sum_{m'} (\widehat{r}_{n-1,m'})^\alpha \cdot \widehat{\rho}_{m'}^*{}^1} \right) (\widehat{r}_{n-1,m})^\alpha, \end{aligned}$$

where $\widehat{\rho}_m^*{}^0$ and $\widehat{\rho}_m^*{}^1$ are defined in Appendix A. Let $m, m' \in K$ with $m \neq m'$ and $r_{1,m} > r_{1,m'} > 0$. Fix $n > 1$, we first show that the rich-get-richer dynamic applies to the ranking probabilities. To simplify notation, let $x \equiv \widehat{r}_{n-1,m}$ and $y \equiv \widehat{r}_{n-1,m'}$. Because $0 < \nu < 1$ we always have $x > y > 0$ for any $n > 2$, and because the two websites have the same signal they also have the equal coefficients on $(\widehat{r}_{n-1,m})^\alpha$ and $\widehat{r}_{n-1,m}$, say, a and b respectively, where $a, b > 0$. Hence we can write:

$$\widehat{r}_{n,m} = ax^\alpha + bx \quad \text{and} \quad \widehat{r}_{n,m'} = ay^\alpha + by.$$

But then it follows that:

$$\frac{\widehat{r}_{n,m}}{\widehat{r}_{n,m'}} > \frac{\widehat{r}_{n-1,m}}{\widehat{r}_{n-1,m'}} \iff \frac{ax^\alpha + bx}{ay^\alpha + by} > \frac{x}{y} \iff ax^\alpha y + bxy > axy^\alpha + bxy \iff \left(\frac{x}{y}\right)^\alpha > \frac{x}{y} \iff \alpha > 1$$

and similarly

$$\frac{\widehat{r}_{n,m}}{\widehat{r}_{n,m'}} \stackrel{(&=)}{<} \frac{\widehat{r}_{n-1,m}}{\widehat{r}_{n-1,m'}} \iff \frac{ax^\alpha + bx}{ay^\alpha + by} \stackrel{(&=)}{<} \frac{x}{y} \iff ax^\alpha y + bxy \stackrel{(&=)}{<} axy^\alpha + bxy \iff \left(\frac{x}{y}\right)^\alpha \stackrel{(&=)}{<} \frac{x}{y} \iff \alpha \stackrel{(&=)}{<} 1$$

Finally, to see the claim, notice that $\hat{\rho}_{n,m} = \frac{a}{1-\nu}(\hat{r}_{n,m})^\alpha$ and $\hat{\rho}_{n,m'} = \frac{a}{1-\nu}(\hat{r}_{n,m'})^\alpha$, where again $\frac{a}{1-\nu} > 0$, so that:

$$\begin{aligned} \frac{\hat{\rho}_{n,m}}{\hat{\rho}_{n,m'}} > \frac{\hat{\rho}_{n-1,m}}{\hat{\rho}_{n-1,m'}} &\iff \frac{\frac{a}{1-\nu}(\hat{r}_{n,m})^\alpha}{\frac{a}{1-\nu}(\hat{r}_{n,m'})^\alpha} > \frac{\frac{a}{1-\nu}(\hat{r}_{n-1,m})^\alpha}{\frac{a}{1-\nu}(\hat{r}_{n-1,m'})^\alpha} \iff \left(\frac{\hat{r}_{n,m}}{\hat{r}_{n,m'}}\right)^\alpha > \left(\frac{\hat{r}_{n-1,m}}{\hat{r}_{n-1,m'}}\right)^\alpha \iff \frac{\hat{r}_{n,m}}{\hat{r}_{n,m'}} > \frac{\hat{r}_{n-1,m}}{\hat{r}_{n-1,m'}} \\ &\iff \alpha > 1, \end{aligned}$$

and correspondingly for $\alpha < 1$ and $\alpha = 1$. \square

Proof of Corollary 1. This follows directly from Proposition 1. Fix a generic initial ranking and consider the website carrying the website majority signal that is highest ranked in the initial ranking, say $\ell_0 \in K$, and the website carrying the website minority signal that is highest ranked in the initial ranking, say ℓ_1 . Because $r_{1,m} > 0$, for all m , and $\gamma > 0$, it is easy to see from Equation (3) that $\rho_{n,m} > 0$ for all n, m . From Proposition 1, we have that, for $\alpha > 1$, $\frac{\rho_{n,\ell_0}}{\rho_{n,m}} > \frac{\rho_{1,\ell_0}}{\rho_{1,m}}$ for any $n > 1$ and $m \in K$ with $m \neq \ell_0$, and also $\frac{\rho_{n,\ell_1}}{\rho_{n,m'}} > \frac{\rho_{1,\ell_1}}{\rho_{1,m'}}$ for any $n > 1$ and $m' \in M \setminus K$ with $m' \neq \ell_1$. In particular, for given k and ζ , by choosing N sufficiently large and $\ell, \ell' \in M'$ we can satisfy $M' = k > 1^{33}$ and $\rho_{n,M'} > \zeta$ for all $n \geq N$, and it immediately follows from $\rho_{n,M \setminus M'} = 1 - \rho_{n,M'}$ that $\rho_{n,M \setminus M'} < 1 - \zeta$. To see how this can be done with a generic initial ranking arbitrarily close to the uniform distribution, or given $\epsilon > 0$ arbitrarily small, repeat the above argument to a generic initial ranking r_1 also satisfying $\sum_n (r_{1,m} - \frac{1}{M})^2 < \epsilon$. This completes the proof. \square

Proof of Proposition 2. We need to check two cases: the websites in L carry the majority signal and the websites in L do not carry the majority signal. In the first case, $\hat{\rho}_{N,K} \geq \hat{\rho}_{N,K'}$ follows from Proposition 4. To see $\hat{\rho}_{N,M \setminus K'} \geq \hat{\rho}_{N,M \setminus K}$ in this case, notice that

$$\hat{\rho}_{N,M \setminus K'} = 1 - \hat{\rho}_{N,K'} \geq 1 - \hat{\rho}_{N,K} = \hat{\rho}_{N,M \setminus K}.$$

Similarly, in the second case, $\hat{\rho}_{N,M \setminus K'} \geq \hat{\rho}_{N,M \setminus K}$ follows from Proposition 4, and $\hat{\rho}_{N,K} \geq \hat{\rho}_{N,K'}$ follows as just above. This concludes the proof. \square

Proof of Proposition 3. Recall the equation defining the ranking probabilities in the case of personalization (Equations (6) and (7)), for any t and m , and for $\ell = A, B$:

$$r_{t,m}^\ell = \nu_t^\ell r_{t-1,m} + (1 - \nu_t^\ell) \rho_{t-1,m}^\ell,$$

where:

$$(\nu_t^\ell, 1 - \nu_t^\ell) = \left(\frac{\kappa}{1 - \lambda_t + \kappa}, \frac{1 - \lambda_t}{1 - \lambda_t + \kappa} \right) \text{ and where } \lambda_t = \begin{cases} 0 & \text{if } t-1 \in \ell \\ \lambda & \text{else.} \end{cases}$$

We can apply the mean dynamics approximation and obtain the deterministic recursions:

$$\begin{aligned} \hat{r}_{t,m}^A &= \hat{r}_{t-1,m}^A + \frac{1}{1 + \kappa_t} (\hat{\rho}_{t-1,m}^A - \hat{r}_{t-1,m}^A) + \frac{1 - \lambda}{1 - \lambda + \kappa_t} (\hat{\rho}_{t-1,m}^B - \hat{r}_{t-1,m}^A) \\ \hat{r}_{t,m}^B &= \hat{r}_{t-1,m}^B + \frac{1}{1 + \kappa_t} (\hat{\rho}_{t-1,m}^B - \hat{r}_{t-1,m}^B) + \frac{1 - \lambda}{1 - \lambda + \kappa_t} (\hat{\rho}_{t-1,m}^A - \hat{r}_{t-1,m}^B), \end{aligned}$$

³³It is easy to see that by choosing γ sufficiently small, the probability of choosing websites carrying the minority signal can be made arbitrarily small and the the probability of choosing websites in K can be made arbitrarily large, so that, in particular, one can choose $k = 1$.

where, following Equation (14), we can write, for $\ell = A, B$:

$$\widehat{\rho}_{t-1,m}^\ell = \mathbb{E}[\rho_{t-1,m}^\ell] = \frac{p_\ell \cdot (\widehat{r}_{t-1,m}^\ell)^\alpha \cdot \widehat{\rho}_m^{\ast 0}}{\sum_{m'} (\widehat{r}_{t-1,m'}^\ell)^\alpha \cdot \widehat{\rho}_{m'}^{\ast 0}} + \frac{(1-p_\ell) \cdot (\widehat{r}_{t-1,m}^\ell)^\alpha \cdot \widehat{\rho}_m^{\ast 1}}{\sum_{m'} (\widehat{r}_{t-1,m'}^\ell)^\alpha \cdot \widehat{\rho}_{m'}^{\ast 1}}.$$

Replacing \widehat{r}_{t-1}^A with $x = (x_1, \dots, x_M)$ and replacing \widehat{r}_{t-1}^B with $y = (y_1, \dots, y_M)$, we can study the function $g : \Delta_\epsilon(M) \times \Delta_\epsilon(M) \rightarrow \mathbb{R}^{2M}$, defined, for $\ell = A, B$, $m = 1, \dots, M$, by:

$$g_m^A(x, y) = \theta_m^A(x, y) - x_m \text{ and } g_m^B(x, y) = \theta_m^B(x, y) - y_m$$

where $\theta^\ell : \Delta_\epsilon(M) \times \Delta_\epsilon(M) \rightarrow \mathbb{R}^M$, $\ell = A, B$, is defined by,

$$\begin{aligned} \theta_m^A(x, y) &= \frac{1}{2-\lambda} \left(\frac{p_A \cdot (x_m)^\alpha \cdot \widehat{\rho}_m^{\ast 0}}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{\rho}_{m'}^{\ast 0}} + \frac{(1-p_A) \cdot (x_m)^\alpha \cdot \widehat{\rho}_m^{\ast 1}}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{\rho}_{m'}^{\ast 1}} \right) + \frac{1-\lambda}{2-\lambda} \left(\frac{p_B \cdot (y_m)^\alpha \cdot \widehat{\rho}_m^{\ast 0}}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{\rho}_{m'}^{\ast 0}} + \frac{(1-p_B) \cdot (y_m)^\alpha \cdot \widehat{\rho}_m^{\ast 1}}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{\rho}_{m'}^{\ast 1}} \right) \\ \theta_m^B(x, y) &= \frac{1-\lambda}{2-\lambda} \left(\frac{p_A \cdot (x_m)^\alpha \cdot \widehat{\rho}_m^{\ast 0}}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{\rho}_{m'}^{\ast 0}} + \frac{(1-p_A) \cdot (x_m)^\alpha \cdot \widehat{\rho}_m^{\ast 1}}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{\rho}_{m'}^{\ast 1}} \right) + \frac{1}{2-\lambda} \left(\frac{p_B \cdot (y_m)^\alpha \cdot \widehat{\rho}_m^{\ast 0}}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{\rho}_{m'}^{\ast 0}} + \frac{(1-p_B) \cdot (y_m)^\alpha \cdot \widehat{\rho}_m^{\ast 1}}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{\rho}_{m'}^{\ast 1}} \right). \end{aligned}$$

Given that the function g is smooth in x, y on $\Delta_\epsilon(M)^2$, it can be shown that the expected limit of our stochastic process can be obtained by solving the system of ordinary differential equations $(\dot{x}, \dot{y}) = g(x, y)$ (or $(\dot{x}, \dot{y}) = (g^A(x, y), g^B(x, y))$) as done above in the case of $\lambda = 0$. We have that when $\lambda = 0$, then it is as if there were a single group with the same (common) ranking algorithm and where the individuals' accuracy is $p = \frac{p_A + p_B}{2}$. As λ increases, then the rankings of the two groups drift apart and it is as if group A had accuracy $\frac{1-\lambda}{2-\lambda}p_A + \frac{1-\lambda}{2-\lambda}p_B$ and group B had accuracy $\frac{1-\lambda}{2-\lambda}p_A + \frac{1}{2-\lambda}p_B$ until it is as if there were two separate rankings of two groups with accuracy p_A and p_B respectively. Since $p_A \neq p_B$ the difference in the accuracy levels of the two groups $(\frac{\lambda}{2-\lambda}(p_A - p_B))$ is increasing in λ . Finally, since clicking and ranking probabilities coincide in the limit, this translates to increasingly different probabilities of clicking on any given majority website in the two groups and hence, given the definition of \mathcal{BP} , also to a measure $\mathcal{BP}(\mathcal{E}_\lambda)$ that is increasing in λ . \square

Proof of Proposition 4. We begin by characterizing the limit ranking probabilities using the differential equation from Appendix A. Since the initial ranking is uniform, we have that the expected ranking probabilities are equal for websites with the same signal, that is, $\widehat{r}_{n,m} = \widehat{r}_{n,m'}$ for any two websites m, m' with $y_m = y_{m'}$. To simplify notation, let $x = \widehat{r}_{n,L} = L \cdot \widehat{r}_{n,m}$ for $m \in L$ be the total expected ranking probability for the websites in L , and let $x' = 1 - x$ be the total expected ranking probability for the remaining websites in $M \setminus L$. Suppose $0 < L < \frac{M-1}{2}$, then we have:

$$\widehat{\rho}_m^{\ast 0} = \begin{cases} \frac{\gamma}{L} & \text{if } m \in L \\ \frac{1-\gamma}{M-L} & \text{else} \end{cases} \quad \text{and} \quad \widehat{\rho}_m^{\ast 1} = \begin{cases} 0 & \text{if } m \in L \\ \frac{1}{M-L} & \text{else,} \end{cases}$$

and the equations defining the limit probabilities can be reduced to a single equation in x :

$$\theta_L^{\text{minority}}(x; \alpha, \gamma, p) \equiv \frac{p\gamma \left(\frac{x}{L}\right)^\alpha}{\gamma \left(\frac{x}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-x}{M-L}\right)^\alpha} = x. \quad (15)$$

Suppose now $\frac{M}{2} \leq L < M$, then we have:

$$\widehat{\rho}_m^{\ast 0} = \begin{cases} \frac{1}{L} & \text{if } m \in L \\ 0 & \text{else} \end{cases} \quad \text{and} \quad \widehat{\rho}_m^{\ast 1} = \begin{cases} \frac{1-\gamma}{L} & \text{if } m \in L \\ \frac{\gamma}{M-L} & \text{else,} \end{cases}$$

and the equations defining the limit probabilities can be reduced to a single equation in x :

$$\theta_L^{majority}(x; \alpha, \gamma, p) \equiv p + \frac{(1-p)(1-\gamma)\left(\frac{x}{L}\right)^\alpha}{(1-\gamma)\left(\frac{x}{L}\right)^\alpha + \gamma\left(\frac{1-x}{M-L}\right)^\alpha} = x. \quad (16)$$

To see 1., notice that

$$0 \leq \theta_L^{minority}(x; \alpha, \gamma, p) \leq p \leq \theta_L^{majority}(x; \alpha, \gamma, p) \leq 1,$$

which in particular implies that, for all values of $0 < L < \frac{M}{2} - 1$ (when websites in L have a minority signal; for simplicity, we consider the case where M is even), we have $x \leq p$, and that, for all values of $\frac{M}{2} \leq L < M$ (when websites in L have a majority signal), we have $x \geq p$.

To see 2., we show that the solutions to the above two equations $\theta_L^{minority}(x; \alpha, \gamma, p) = x$ and $\theta_L^{majority}(x; \alpha, \gamma, p) = x$ are decreasing in L in the corresponding ranges. These are determined respectively as the intersections of the functions $\theta_L^{minority}$ and $\theta_L^{majority}$ with the function x . It can be checked that $0 = \theta_L^{minority}(0; \alpha, \gamma, p) \leq \theta_L^{minority}(\frac{1}{2}; \alpha, \gamma, p) \leq p$, for $0 < L < \frac{M-1}{2}$, and $p \leq \theta_L^{majority}(\frac{1}{2}; \alpha, \gamma, p) \leq \theta_L^{majority}(1; \alpha, \gamma, p) = 1$, for $\frac{M}{2} \leq L < M$. In particular, it can be checked that at the relevant solutions ($x_L^{majority}$ and $x_L^{minority}$) the functions $\theta_L^{minority}$ and $\theta_L^{majority}$ intersect the x function with positive a slope that is less than one. Hence, to determine the sign of the effect of L on interim efficiency it suffices to look at the derivatives of $\theta_L^{minority}$ and $\theta_L^{majority}$ with respect to L :

$$\begin{aligned} \frac{\partial \theta_L^{minority}}{\partial L} &= -\frac{\alpha\gamma(1-\gamma)pM\left(\frac{x}{L}\right)^\alpha\left(\frac{1-x}{M-L}\right)^\alpha}{L(M-L)\left(\left(\frac{1-x}{M-L}\right)^\alpha + \gamma\left(\left(\frac{x}{L}\right)^\alpha - \left(\frac{1-x}{M-L}\right)^\alpha\right)\right)^2} \\ \frac{\partial \theta_L^{majority}}{\partial L} &= -\frac{\alpha\gamma(1-\gamma)(1-p)M\left(\frac{x}{L}\right)^\alpha\left(\frac{1-x}{M-L}\right)^\alpha}{L(M-L)\left(\left(\frac{x}{L}\right)^\alpha + \gamma\left(\left(\frac{1-x}{M-L}\right)^\alpha - \left(\frac{x}{L}\right)^\alpha\right)\right)^2}, \end{aligned}$$

which are both clearly non-positive. This further implies that the solutions will satisfy $\frac{\partial x}{\partial L} \leq 0$ and hence $\frac{\partial \hat{\rho}_{N,L}}{\partial L} \leq 0$ on the relevant ranges and for N sufficiently large. The statement then follows immediately from the definition of \mathcal{P}_L . \square

Proof of Proposition 5. Notice that from the definition of \mathcal{P} in Equation (12), we immediately get:

$$\frac{\partial \mathcal{P}}{\partial z} = \sum_{L=0}^M \binom{M}{L} q^L (1-q)^{M-L} \frac{\partial \mathcal{P}_L}{\partial z},$$

for any variable z . Hence, we can evaluate the comparative statics by looking at the effects on the interim efficiency, $(\frac{\partial \mathcal{P}_L}{\partial z})$. So to see that \mathcal{P} is weakly increasing in p , since the initial ranking is uniform, we can use the same reasoning as in the proof of Proposition 4. In particular, it suffices to consider the following derivatives for $\theta_L^{minority}$ and $\theta_L^{majority}$ defined respectively in Equations (15) and (16) above:

$$\frac{\partial \theta_L^{minority}}{\partial p} = \frac{\gamma\left(\frac{x}{L}\right)^\alpha}{\gamma\left(\frac{x}{L}\right)^\alpha + (1-\gamma)\left(\frac{1-x}{M-L}\right)^\alpha} \geq 0, \quad \frac{\partial \theta_L^{majority}}{\partial p} = 1 - \frac{(1-\gamma)\left(\frac{x}{L}\right)^\alpha}{\gamma\left(\frac{1-x}{M-L}\right)^\alpha + (1-\gamma)\left(\frac{x}{L}\right)^\alpha} \geq 0.$$

This implies that $\rho_{\infty,L} = \mathcal{P}_L$ is weakly increasing in p for all values of L and hence so is \mathcal{P} . To see

that, \mathcal{P} is weakly decreasing in γ , notice that:

$$\frac{\partial \theta_L^{\text{minority}}}{\partial \gamma} = \frac{p \left(\frac{x}{L}\right)^\alpha \left(\frac{1-x}{M-L}\right)^\alpha}{\left(\gamma \left(\frac{x}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-x}{M-L}\right)^\alpha\right)^2} \geq 0, \quad \frac{\partial \theta_L^{\text{majority}}}{\partial \gamma} = -\frac{(1-p) \left(\frac{x}{L}\right)^\alpha \left(\frac{1-x}{M-L}\right)^\alpha}{\left((1-\gamma) \left(\frac{x}{L}\right)^\alpha + \gamma \left(\frac{1-x}{M-L}\right)^\alpha\right)^2} \leq 0.$$

It can be further checked that for $q > \frac{1}{2}$, the negative effect of when the websites in L have a majority signal outweighs the effect of when they are a minority signal. We here show it explicitly for the two values $\gamma = 0$ and $\gamma = 1$. Notice that the solutions to the Equations (15) and (16) of the proof to Proposition 4 are, respectively, $x_L^{\text{minority}} = 0$ and $x_L^{\text{majority}} = 1$ for $\gamma = 0$ and are, respectively, $x_L^{\text{minority}} = p$ and $x_L^{\text{majority}} = p$ for $\gamma = 1$. (Recall that $x_L^{\text{minority}} = \hat{\rho}_{\infty, L}$ when L is a website-minority website and $x_L^{\text{majority}} = \hat{\rho}_{\infty, L}$ when L is a website-majority website.) For the derivatives, this implies:

$$\left. \frac{\partial \theta_L^{\text{minority}}}{\partial \gamma} \right|_{\gamma=0} = 0, \quad \left. \frac{\partial \theta_L^{\text{majority}}}{\partial \gamma} \right|_{\gamma=0} = 0 \quad \text{and} \quad \left. \frac{\partial \theta_L^{\text{minority}}}{\partial \gamma} \right|_{\gamma=1} = p, \quad \left. \frac{\partial \theta_L^{\text{majority}}}{\partial \gamma} \right|_{\gamma=1} = -(1-p).$$

For the case $\gamma = 0$ we are done. For the case $\gamma = 1$, we have for the derivative:

$$\sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} q^L (1-q)^{M-L} p - \sum_{L=\frac{M}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} (1-p) < 0,$$

immediately follows from $\frac{1}{2} < p \ll q$. This in turn implies that $\left. \frac{\partial \mathcal{P}}{\partial \gamma} \right|_{\gamma=1} < 0$.

We now move to q and consider first the cases $\gamma = 0, \gamma = 1$. As mentioned above, when $\gamma = 0$, we have that $x_L^{\text{minority}} = 0$ and $x_L^{\text{majority}} = 1$, whereas for $\gamma = 1$, we have that $x_L^{\text{minority}} = p$ and $x_L^{\text{majority}} = p$ for $\gamma = 1$. As a result, we have (again, we consider for simplicity the case where M is even):

$$\mathcal{P}|_{\gamma=0} = \sum_{L=\frac{M}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} + q^M \quad \text{and} \quad \mathcal{P}|_{\gamma=1} = \sum_{L=1}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} p + q^M,$$

which are strictly increasing in q , since a higher q shifts probability towards states with strictly higher interim efficiency (either from 0 to 1 when $\gamma = 0$ or from p to 1 when $\gamma = 1$). Because the functions $\theta_L^{\text{minority}}, \theta_L^{\text{majority}}$ are continuously differentiable in γ , and \mathcal{P} is continuously differentiable in γ and q this means that a constant $\bar{\gamma} > 0$ can be found such that \mathcal{P} is weakly increasing on $[0, \bar{\gamma}] \cup [1 - \bar{\gamma}, 1]$. \mathcal{P} can be decreasing in q for intermediate values of γ due to the fact that a higher q shifts probability towards states with higher levels of L . Hence, given the non-monotonicity of interim efficiency in L for intermediate levels of γ , this may overall strictly decrease \mathcal{P} . To see the effect of α , notice that:

$$\frac{\partial \theta_L^{\text{minority}}}{\partial \alpha} = \frac{\gamma(1-\gamma)p \left(\frac{x}{L}\right)^\alpha \left(\frac{1-x}{M-L}\right)^\alpha \left(\log \frac{x}{L} - \log \frac{1-x}{M-L}\right)}{\left(\gamma \left(\frac{x}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-x}{M-L}\right)^\alpha\right)^2}$$

and

$$\frac{\partial \theta_L^{majority}}{\partial \alpha} = \frac{\gamma(1-\gamma)(1-p) \left(\frac{x}{L}\right)^\alpha \left(\frac{1-x}{M-L}\right)^\alpha \left(\log \frac{x}{L} - \log \frac{1-x}{M-L}\right)}{\left((1-\gamma) \left(\frac{x}{L}\right)^\alpha + \gamma \left(\frac{1-x}{M-L}\right)^\alpha\right)^2}.$$

Consider first the case $\gamma \approx 1$. We have that $x_L^{minority} \approx p$ and $x_L^{majority} \approx p$ and hence we can approximate $\frac{\partial \mathcal{P}}{\partial \alpha}$ by:

$$\begin{aligned} & \gamma(1-\gamma) \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} q^L (1-q)^{M-L} p \frac{\left(\frac{p}{L}\right)^\alpha \left(\frac{1-p}{M-L}\right)^\alpha \left(\log \frac{p}{L} - \log \frac{1-p}{M-L}\right)}{\left(\gamma \left(\frac{p}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{M-L}\right)^\alpha\right)^2} \\ + & \gamma(1-\gamma) \sum_{L=\frac{M}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} (1-p) \frac{\left(\frac{p}{L}\right)^\alpha \left(\frac{1-p}{M-L}\right)^\alpha \left(\log \frac{p}{L} - \log \frac{1-p}{M-L}\right)}{\left((1-\gamma) \left(\frac{p}{L}\right)^\alpha + \gamma \left(\frac{1-p}{M-L}\right)^\alpha\right)^2} \\ = & \gamma(1-\gamma) \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{p^{1+\alpha} (1-p)^\alpha \log \frac{p}{1-p} \log \frac{M-L}{L}}{(L(M-L))^\alpha \left(\gamma \left(\frac{p}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{M-L}\right)^\alpha\right)^2} \\ + & \gamma(1-\gamma) \sum_{L=1}^{\frac{M}{2}} \binom{M}{L} q^{M-L} (1-q)^L \frac{p^\alpha (1-p)^{1+\alpha} \log \frac{p}{1-p} \log \frac{L}{M-L}}{(L(M-L))^\alpha \left(\gamma \left(\frac{p}{M-L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{L}\right)^\alpha\right)^2} \\ = & \gamma(1-\gamma) \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{p^{1+\alpha} (1-p)^\alpha \log \frac{p}{1-p} \log \frac{M-L}{L}}{(L(M-L))^\alpha \left(\gamma \left(\frac{p}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{M-L}\right)^\alpha\right)^2} \\ - & \gamma(1-\gamma) \sum_{L=1}^{\frac{M}{2}} \binom{M}{L} q^{M-L} (1-q)^L \frac{p^\alpha (1-p)^{1+\alpha} \log \frac{p}{1-p} \log \frac{M-L}{L}}{(L(M-L))^\alpha \left(\gamma \left(\frac{p}{M-L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{L}\right)^\alpha\right)^2} \\ \leq & \gamma(1-\gamma) \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} \frac{(p(1-p))^\alpha \log \frac{p}{1-p} \log \frac{M-L}{L}}{(L(M-L))^\alpha} \times \\ & \times \left(\frac{q^L (1-q)^{M-L} p}{\left(\gamma \left(\frac{p}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{M-L}\right)^\alpha\right)^2} - \frac{q^{M-L} (1-q)^L (1-p)}{\left(\gamma \left(\frac{p}{M-L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{L}\right)^\alpha\right)^2} \right) \\ < & 0. \end{aligned}$$

The weak inequality follows from omitting the case $L = \frac{M}{2}$ which is ≤ 0 . The last inequality follows because $q^L (1-q)^{M-L} p < q^{M-L} (1-q)^L (1-p)$ since $1 > q > \frac{1}{2}$ and because:

$$\gamma \left(\frac{p}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{M-L}\right)^\alpha > \gamma \left(\frac{p}{M-L}\right)^\alpha + (1-\gamma) \left(\frac{1-p}{L}\right)^\alpha,$$

since $\gamma \approx 1$ and $L < \frac{M}{2} < M-L$. Again, because the functions $\theta_L^{minority}$, $\theta_L^{majority}$ are continuously differentiable in γ and α , we can find an appropriate $\bar{\gamma}$ and the claim follows.

Finally, to see that \mathcal{P} can be increasing in α for small values of γ , consider, for example, a search environment with $\gamma = \frac{1}{2}$, $\alpha = 1$ and $q = 0.75$. Then it can be shown that a slight increase in α will strictly increase ex ante efficiency. \square

Proof of Corollary 2 This follows from Proposition 5. It is sufficient to compare \mathcal{P} at $\alpha = 0$ ($\mathcal{P}|_{\alpha=0}$) with \mathcal{P} at a fixed value $\alpha > 0$ ($\mathcal{P}|_{\alpha>0}$). We know that $\mathcal{P}|_{\alpha=0} = q$. Since $M > 1$ and $q > p > \frac{1}{2}$, we also know that $\mathcal{P}|_{\alpha>0} > q$ for $\gamma = 0$ and that $\mathcal{P}|_{\alpha>0} \approx p$ for $\gamma = 1$. Since $q > p$ and $\mathcal{P}|_{\alpha>0}$ is continuous and weakly decreasing in γ , there exists $\bar{\gamma}$ such that $\mathcal{P}|_{\alpha>0} \geq q$ for $\gamma \in [0, \bar{\gamma}]$ and $\mathcal{P}|_{\alpha>0} \leq q$ for $\gamma \in [\bar{\gamma}, 1]$, and the claim follows. \square

Proof of Proposition 6. Let \mathcal{P}^λ and \mathcal{P}_L^λ denote respectively ex ante and interim efficiency as a function of the personalization parameter λ . Because the initial ranking is uniform we can look again at the solutions to the Equations (15) and (16) in the proof of Proposition 4, where we recall again that $x_L^{minority} = \hat{\rho}_{\infty,L}$ when websites in L have minority signal and $x_L^{majority} = \hat{\rho}_{\infty,L}$ when they have majority signal. Hence to show that \mathcal{P}^λ is weakly decreasing in λ we can use these solutions to study the interim efficiency levels \mathcal{P}_L^λ for $\lambda \in [0, 1]$ and $1 < L < M$. We argue that it is sufficient to show that \mathcal{P}^λ is concave in p , essentially since it implies that the efficiency of an average group will be below the average of the efficiency of the two groups. Since \mathcal{P}^λ is a weighted average of the different \mathcal{P}_L^λ 's, it is enough to show that these are on average sufficiently concave in p . Fix a realization of \mathcal{E} , say, parametrized by L . From the proof of Proposition 3, we have that when $\lambda = 0$, it is as if there were a single group with a common ranking algorithm, with $p = \frac{p_A + p_B}{2}$. As λ increases, then the rankings of the two groups drift apart and are as if group A had accuracy $\frac{1-\lambda}{2-\lambda}p_A + \frac{1-\lambda}{2-\lambda}p_B$ and group B had accuracy $\frac{1-\lambda}{2-\lambda}p_A + \frac{1}{2-\lambda}p_B$, until, when $\lambda = 1$, it is as if there were two separate rankings of two groups with accuracy p_A and p_B , respectively. As a result interim efficiency can be written as the average of the interim efficiency of two groups, one with ranking accuracy $\frac{1-\lambda}{2-\lambda}p_A + \frac{1-\lambda}{2-\lambda}p_B$ and clicking with accuracy p_A and another with ranking accuracy $\frac{1-\lambda}{2-\lambda}p_A + \frac{1}{2-\lambda}p_B$ and clicking with accuracy p_B . Therefore, in the case of intermediate values of λ interim efficiency can be seen as the average of two levels of interim efficiency corresponding to two different signal accuracies that are increasingly apart as λ increases (that is, go from both signals corresponding to $\frac{p_A + p_B}{2}$ when $\lambda = 0$ to being p_A and p_B respectively when $\lambda = 1$). Therefore, it suffices to show that the basic interim efficiency function \mathcal{P}_L^λ is concave in p . We show this by evaluating the first and second derivatives of $\hat{\rho}_{\infty,L} = \mathcal{P}_L^\lambda$. We already know by Proposition 5 that the first derivatives with respect to p are positive. We now study the second derivatives. Let

$$H_L^{minority}(x; \alpha, \gamma, p) = \theta_L^{minority}(x; \alpha, \gamma, p) - x \quad \text{and} \quad H_L^{majority}(x; \alpha, \gamma, p) = \theta_L^{majority}(x; \alpha, \gamma, p) - x,$$

then, at the solutions $x = x_L^{minority}$ and $x = x_L^{majority}$, we have that $H_L^{minority}(x_L^{minority}; \alpha, \gamma, p) = 0$ and $H_L^{majority}(x_L^{majority}; \alpha, \gamma, p) = 0$, respectively, and hence the first derivatives of the solutions with respect to p can be written as:

$$\frac{dx_L^{minority}}{dp} = -\frac{\partial H_L^{minority}(x_L^{minority}; \alpha, \gamma, p)/\partial p}{H_L^{minority}(x_L^{minority}; \alpha, \gamma, p)/\partial x} \quad \text{and} \quad \frac{dx_L^{majority}}{dp} = -\frac{\partial H_L^{majority}(x_L^{majority}; \alpha, \gamma, p)/\partial p}{H_L^{majority}(x_L^{majority}; \alpha, \gamma, p)/\partial x}.$$

From this we can compute the second derivatives as:

$$\frac{d^2 x_L^{minority}}{dp^2} = \frac{\alpha \gamma^2 (1 - \gamma) x (1 - x) \left(\frac{x}{L}\right)^{2\alpha} \left(\frac{1-x}{M-L}\right)^\alpha \left(\gamma \left(\frac{x}{L}\right)^\alpha + (1 - \gamma) \left(\frac{1-x}{M-L}\right)^\alpha\right)}{\left(\alpha \gamma (1 - \gamma) p \left(\frac{x}{L}\right)^\alpha \left(\frac{1-x}{M-L}\right)^\alpha - x (1 - x) \left(\gamma \left(\frac{x}{L}\right)^\alpha + (1 - \gamma) \left(\frac{1-x}{M-L}\right)^\alpha\right)^2\right)^2} \geq 0$$

and

$$\frac{d^2 x_L^{majority}}{dp^2} = - \frac{\alpha \gamma^2 (1-\gamma) x(1-x) \left(\frac{x}{L}\right)^\alpha \left(\frac{1-x}{M-L}\right)^{2\alpha} \left((1-\gamma) \left(\frac{x}{L}\right)^\alpha + \gamma \left(\frac{1-x}{M-L}\right)^\alpha \right)}{\left(\alpha \gamma (1-\gamma) (1-p) \left(\frac{x}{L}\right)^\alpha \left(\frac{1-x}{M-L}\right)^\alpha - x(1-x) \left((1-\gamma) \left(\frac{x}{L}\right)^\alpha + \gamma \left(\frac{1-x}{M-L}\right)^\alpha \right)^2 \right)^2} \leq 0.$$

This shows that the solutions $x_L^{minority}$ and hence the \mathcal{P}_L^λ when L are minority websites are convex in p , while the solutions $x_L^{majority}$ and hence the \mathcal{P}_L^λ when L are majority websites are concave in p . However, given the linearity of the differential operator, the second derivative of the overall ex ante efficiency function is given by the weighted sum:

$$\frac{d^2 \mathcal{P}^\lambda}{dp^2} = \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{d^2 x_L^{minority}}{dp^2} + \sum_{L=\frac{M}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{d^2 x_L^{majority}}{dp^2}$$

evaluated at the corresponding solutions (again, we consider for simplicity the case where M is even). Now, since $p \leq \bar{p}$, for $\bar{p} > \frac{1}{2}$ not too large, it can be shown that if $x_L^{minority}$ is a solution to $H_L^{minority}(x; \alpha, \gamma, p) = 0$ for $L \leq \frac{M}{2} - 1$, then the solution $x_{M-L}^{majority}$ to $H_{M-L}^{majority}(x; \alpha, \gamma, p) = 0$,³⁴ is arbitrarily close to $1 - x_L^{minority}$ for the given L . As a result, we have that, for $1 < L \leq \frac{M}{2} - 1$,

$$\left. \frac{d^2 x_{M-L}^{majority}}{dp^2} \right|_{x=x_{M-L}^{majority}} \approx \left. \frac{d^2 x_{M-L}^{majority}}{dp^2} \right|_{x=1-x_L^{minority}} \approx - \left. \frac{d^2 x_L^{minority}}{dp^2} \right|_{x=x_L^{minority}}$$

and hence we can write:

$$\begin{aligned} \frac{d^2 \mathcal{P}^\lambda}{dp^2} &= \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{d^2 x_L^{minority}}{dp^2} + \sum_{L=\frac{M}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{d^2 x_L^{majority}}{dp^2} \\ &\leq \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{d^2 x_L^{minority}}{dp^2} + \sum_{L=\frac{M}{2}+1}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{d^2 x_L^{majority}}{dp^2} \\ &= \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{d^2 x_L^{minority}}{dp^2} - \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} q^{M-L} (1-q)^L \frac{d^2 x_L^{minority}}{dp^2} \\ &= \sum_{L=1}^{\frac{M}{2}-1} \binom{M}{L} \frac{d^2 x_L^{minority}}{dp^2} (q^L (1-q)^{M-L} - q^{M-L} (1-q)^L) \leq 0. \end{aligned}$$

Again, the weak inequality follows from omitting the case $L = \frac{M}{2}$ which is anyways concave in p . The strict inequality follows since $q > p_A, p_B$. This shows that the weighted sum that yields \mathcal{P}^λ is overall concave in p which completes the proof. \square

³⁴Where by slight abuse of notation $M - L$ here corresponds to a majority of outlets with correct signal.