



**Universitat
Pompeu Fabra**
Barcelona

Department
of Economics and Business

**Economic Working Paper Series
Working Paper No. 1552**

**Opinion Dynamics via search engines
(and other algorithmic gatekeepers)**

Fabrizio Germano and Francesco Sobbrío

Updated version: March 2018
(December 2016)

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

Fabrizio Germano[†] and Francesco Sobbrío[‡]

First Version: December 2016

This Version: March 2018

Abstract

Ranking algorithms are the information gatekeepers of the Internet era. We develop a stylized model to study the effects of ranking algorithms on opinion dynamics. We consider a search engine using an algorithm that depends on popularity and on personalization. Popularity-based rankings generate an advantage of the fewer effect: fewer websites reporting a given signal attract more traffic overall. This provides a rationale for the diffusion of misinformation, as traffic to websites reporting incorrect information can be large precisely when there are few of them. Finally, we study conditions under which popularity-based rankings and personalized rankings contribute to asymptotic learning.

Keywords: Ranking Algorithms, Opinion Dynamics, Website Traffic, Asymptotic Learning, Stochastic Choice, Misinformation, Polarization, Search Engines, Fake News.

JEL Classification: D83, L86

*We thank Larbi Alaoui, Jose Apesteguia, Emilio Calvano, Stefano Colombo, Lisa George, Roberto Imbuzeiro Oliveira, Matthieu Manant, Andrea Mattozzi, Ignacio Monzón, Antonio Nicolò, Nicola Persico, Vaiva Petrikaite, Christian Peukert, Alessandro Riboni, 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), Rome (IMEBESS 2016) and Toulouse (Digital Economics Conference 2018) for useful comments and conversations. We are also indebted to Riccardo Boscolo for helpful discussions on the functioning of search engines. Germano 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”) and also from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563).

[†]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”, and CESifo. 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 them 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; see *Google Trends*). Remarkably, what determines the ranking of any website to be displayed for any given search query are automated algorithms. Such algorithms are also used by social media—such as Facebook and Twitter—to rank their posts or tweets. They are so fundamental in establishing what is relevant information or what are relevant information sources, that they necessarily convert search engines and social media platforms into *de facto* “algorithmic gatekeepers” (Introna and Nissenbaum, 2000; Rieder, 2005; Granka, 2010; Napoli, 2015; Tufekci, 2015).

Despite the importance of such ranking algorithms for a wide variety of online platforms, opinion dynamics via algorithmic rankings is largely understudied. This paper aims to fill this gap by developing a dynamic framework that studies the interaction between individual searches and a stylized ranking algorithm. The search engine provides a ranking of the available websites that is used by the individuals to look for information on the state of the world (e.g., whether or not to vaccinate a child). Individuals’ choices are modeled by means of a random utility model (Luce, 1959; Block and Marschak, 1960; Gül *et al.*, 2014) that yields probabilities of reading the ranked websites that depend on both website content and ranking. Individuals thus tend to choose websites that are: (i) higher-ranked (De Cornière and Taylor, 2014; Taylor, 2013; Hagiu and Jullien, 2014; Burguet *et al.*, 2015), (ii) *ex-ante* more informative, (iii) confirm their prior information (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2010), and (iv) they may trade off content and ranking as a result of their stochastic choices. Website rankings are endogenous meaning that individuals’ choices feed back into the search engine’s rankings, thereby affecting future searches (Demange, 2014a). We focus on two key aspects of ranking algorithms: (a) rankings may be based on the *popularity* of the different websites, (b) rankings may be *personalized* and may depend on individuals’ characteristics.

The model provides four main insights deriving from the interaction of the ranking algorithm with sequential individual searches. First, we uncover a fundamental property induced by *popularity-based* rankings, which we call the *advantage of the fewer (AOF)*. It says that fewer websites carrying a given signal may attract more traffic overall, than if there were more of them.

Popularity-based rankings amplify the static effect of simply concentrating audiences on fewer outlets, since they induce relatively higher rankings for the fewer outlets, which makes them more attractive for subsequent individuals, further raising their ranking and so on, generating a “*few get richer*” dynamic with a potentially sizeable effect in the limit. We show that *AOF* holds even beyond the assumptions of our basic model. Importantly, it may potentially provide a rationale to explain why minority websites (e.g., “conspiracy” or “fake news” websites) may thrive and gain in authority in the current information environment, dominated by algorithmic gatekeepers, such as search engines and social media.¹

Second, we provide conditions under which *popularity-based* rankings can effectively aggregate information. We compare asymptotic learning under popularity-based rankings and under fully random rankings. We find that popularity-based rankings do better as long as individuals provide sufficiently positive feedback through their searches (e.g., their private signals are not too noisy, or their preference for confirmatory news are not too strong).

Third, we study whether *personalized* rankings can contribute to information aggregation. We compare asymptotic learning under both personalized and non-personalized rankings. We find that there is a close relationship between the conditions under which non-personalized rankings outperform personalized rankings with the ones under which popularity-based rankings outperform random rankings. For the common-value searches of our model, personalization limits the feedback among individuals in the opinion dynamic, making it better *not* to personalize the ranking when individuals’ choices over websites generate a positive feedback. As a consequence, personalized rankings are often dominated in terms of asymptotic learning either by non-personalized rankings or by random rankings. Finally, we also show that personalized rankings can induce relatively similar individuals to read different websites, which can lead to *belief polarization*.

We conclude with some caveats. First, our analysis is positive and we do not pursue a mechanism design approach.² Our aim is to isolate two important components of ranking algorithms—popularity and personalization—and to study their interplay with individuals’ search behavior. Second, while our results suggest that personalization may often be sub-optimal in the context of individuals looking for information on common value issues (e.g., whether or not to vaccinate a child), when individuals care differently about the objects of their searches (e.g., where to have dinner) an appropriately personalized search algorithm might clearly outperform a non-personalized

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).

²Kremer *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.

one. Finally, we are aware that the actual algorithms used by platforms such as Google or Facebook are highly complex.³ As a result, we do not aim to pin down the exact pattern of website traffic or opinion dynamic generated by any specific algorithm.

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 search engines.⁴

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).

Our focus on 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, where 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 automated algorithms (Granka, 2010; Tufekci, 2015).⁵ Therefore, whatever *bias* might originate from search engines, its nature is *intrinsically* different from one arising in, say, traditional news media. As a result, studying the effects of search engines on the accuracy of individuals’ beliefs, requires a different approach from the ones

³As noted by Lazer (2015), such complexity is amplified by the interaction of the algorithm itself with individual search behavior (“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). This suggests that a theoretical model studying the interplay between users’ search behavior and fundamental aspects of the ranking algorithm may yield insights that may otherwise be difficult to identify simply looking at the algorithm code or at observational data.

⁴The 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; Hagiü and Jullien, 2014; Burguet *et al.*, 2015). See also Grimmelmänn (2009) and Hazan (2013) for a legal perspective on the issue.

⁵Put 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.

used so far 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 (or, more generally, of algorithmic gatekeeping) on democratic outcomes. This literature encompasses communication scholars (Hargittai, 2004; Granka, 2010), legal scholars (Goldman, 2006; Grimmelman, 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 presents our central result concerning the effects of the search engine's ranking algorithm on website traffic: the advantage of the fewer. Section 4 presents the implications of the model in terms of asymptotic learning, also providing a comparison with a fully randomized ranking. Section 5 discusses the effects of personalization of search results on belief polarization and asymptotic learning. Section 6 discusses some extensions which assess the robustness of the results of the benchmark model and provide further insights. Section 7 concludes. Finally, all the proofs and 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 (e.g., whether or not to vaccinate a child). Individuals get 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 characterized by its ranking algorithm, which ranks and directs individuals to different websites, using, among other things, the popularity of individuals' choices. To simplify the analysis, we assume that individuals perform exactly one search, one after the other, without knowing who searched before them and without updating prior beliefs after observing the ranking. We 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 the true state ω ,

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

This determines the *website majority signal*, denoted $y_K \in \{0, 1\}$, which is the signal that is carried by a majority of websites; let $K = \{m \in M \mid y_m = y_K\}$ denote the set (and number) of websites carrying the signal y_K . Similarly, each individual $n \in N$ receives two private random signals:

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

which reflects the individual's prior on the true state of the world, and,

$$z_n \in \{0, 1\} \text{ with } \mathbb{P}(z_n = y_K \mid y_K = \zeta) = \mu \in \left(\frac{1}{2}, 1\right], \text{ for any } \zeta \in \{0, 1\},$$

which is independent of x_n , when conditioned on (ω, y_K) , and reflects the individual's prior about what the majority of websites is reporting.

We assume all signals to be conditionally independent across individuals and websites, when conditioned on (ω, y_K) , and that $\mu q > p$.⁶ This very last condition implies not only that, from an ex-ante perspective, the website signals (y_m) and hence also the website majority signal (y_K) are more informative than the individual private signals on the state of the world (x_n), but also that, for a given individual, his private *perceived* website majority signal (z_n) is also more informative than his other signal (x_n). Among other things, the signals z_n allow us to introduce a “rational” side to individuals' choices as we will discuss further below. From the onset, we emphasize that the case $\mu < 1$ is meant to capture the intrinsic noise faced by individuals in identifying the ex-ante more informative website majority signal (y_K).⁷ All conclusions remain unchanged if we assume that individuals can perfectly identify this signal ($\mu = 1$).

⁶This condition implies:

$$\mu \sum_{K > \frac{M}{2}} \binom{M}{K} q^K (1-q)^{M-K} + (1-\mu) \sum_{K > \frac{M}{2}} \binom{M}{K} q^{M-K} (1-q)^K > p \quad \text{and} \quad q > p,$$

where the first inequality implies that an individual's signal z_n is ex ante more informative about ω than the same individual's signal x_n ; and the second inequality implies that a website's signal y_m is ex-ante more informative about ω than an individual's signal x_n .

⁷Accordingly, $(1-\mu)$ represents the probability than an individual incorrectly identifies the minority signal as the majority one (e.g., looking at a website reporting “alternative” facts believing that it is reporting mainstream news).

2.2 Websites

Each of the M websites is characterized by a signal $y_m \in \{0, 1\}$ as described above. Websites report their signal and keep it fixed throughout the search process. They are therefore non-strategic players that carry information that they provide for free to any individual accessing their website.⁸ For simplicity and unless otherwise noted, we assume that M is an odd number.

2.3 Individuals

Individuals in N enter in a random order, sequentially, such that, at any point in time t , there is a unique individual $t \equiv n \in N$, who receives random (i.i.d.) signals $(x_n, z_n) \in \{0, 1\}^2$ as specified above, who performs exactly one search, faces the ranking and chooses a website to read.⁹ When choosing which website to read, individuals make stochastic choices derived from a random utility model (Luce, 1959; Block and Marschak, 1960; Gül *et al.*, 2014; Agranov and Ortoleva, 2017, provide empirical evidence),

$$\rho_{n,m} = \frac{v_{n,m}}{\sum_{m' \in M} v_{n,m'}}, \quad (1)$$

where $\rho_{n,m}$ is the probability individual n clicks on website m , and where $v_{n,m}$ is the value derived from clicking on website m . It is to be interpreted as measuring desirability in the sense of a stochastic preference, which, moreover, reflects in separable form, a ranking aspect and a content aspect of website m . We define the values $v_{n,m}$ in two steps.

Website content. Suppose individuals see all websites ranking-free, that is, with equal probability (fixing content). As mentioned, each individual receives the signals $(x_n, z_n) \in \{0, 1\}^2$. We use these signals to define *attributes* (see Gül *et al.*, 2014) that yield value to the individual in the following sense. Reading a website m , with $y_m = x_n$, yields value $\gamma \in [0, 1]$, where the parameter γ calibrates the preference for reading like-minded news (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2010).¹⁰ At the same time, reading a website m , with $y_m = z_n$, yields value $1 - \gamma$, which reflects the desirability associated with reading the ex-ante most informative signal.¹¹ Using

⁸Besides 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.

⁹For reasons of tractability, we assume that all individuals perform the same search query exactly once.

¹⁰See Yom-Tov *et al.* (2013); Flaxman *et al.* (2013); White and Horvitz (2015) for evidence of individual preference for like-minded news in the context of search engines.

¹¹As mentioned above in Section 2.1, given the assumptions on p, q and μ , a Bayesian rational agent who derives utility from reading a website reporting a correct signal ($y_m = \omega$) but who cannot use the ranking r_n to update her information about the true state of the world, prefers to ignore her signal x_n and choose a website m with $y_m = z_n$. In Section 6.3, we briefly discuss sophisticated learning in our environments.

x_n and z_n to define attributes (with a normalization such that the value of having both attributes is 1), we obtain the *ranking-free values* from selecting website m ,

$$v_{n,m}^* = \begin{cases} \frac{1}{[m]} & \text{if } x_n = y_m = z_n \\ \frac{\gamma}{[m]} & \text{if } x_n = y_m \neq z_n \\ \frac{1-\gamma}{[m]} & \text{if } x_n \neq y_m = z_n \\ 0 & \text{if } x_n \neq y_m \neq z_n \end{cases}, \quad (2)$$

where $[m] = \#\{m' \in M | y_{m'} = y_m\}$ is the number of all websites with the same signal as m . The division by $[m]$ avoids the usual Luce effect when dealing with duplicate alternatives. It also allows us to interpret the $v_{m,n}^*$'s both as values as well as the individual stochastic *ranking-free website choices*. In particular, $v_n^* : \{0, 1\}^2 \rightarrow \Delta(M)$ depends only on the private signals $(x_n, z_n) \in \{0, 1\}^2$.

Website rankings. Let $r_n \in \Delta(M)$ be the *ranking* of the M websites at time $t \equiv n$; the ranking is provided by the search engine (as discussed in the next section). Individuals ignore who searched before them and do not update their beliefs about the true state of the world ω after observing the ranking, but we assume they can process the ranked list, subject to limitations of attention, meaning they may implicitly favor some websites over others, depending on the websites' rankings. Indeed, as shown by Pan *et al.* (2007); Glick *et al.* (2011); Yom-Tov *et al.* (2013); Epstein and Robertson (2015), keeping all other things equal (e.g., the fit of a given website with respect to the individual's preferences), highest ranked websites tend to receive significantly more "attention" by users than lower ranked ones.¹² Since we model the ranking (r_n) as a probability distribution over the outlets that is multiplicative separability from ranking-free values (v_n^*), we can write final values as ranking-weighted values:

$$v_{n,m} = r_{n,m} \cdot v_{n,m}^*. \quad (3)$$

Websites with a higher ranking receive a higher weight.

Individual stochastic choice. We can think of the values $v_{n,m}$ defined in (3) as representing desirability levels for the different ranked alternatives, which yields the stochastic choices defined in (1), as in a standard random utility model, and thus give us the *website choice function* ρ_n :

¹²More generally, when faced with ordinal lists, individuals often show a disproportionate tendency to select options that are placed at the top (Novarese and Wilson, 2013).

$\{0, 1\}^2 \times \Delta(M) \rightarrow \Delta(M)$, defined by:¹³

$$\rho_{n,m} = \frac{r_{n,m} \cdot v_{n,m}^*}{\sum_{m' \in M} r_{n,m'} \cdot v_{n,m'}^*}. \quad (4)$$

The website choices $\rho_{n,m}$ reflect, through the ranking-free values ($v_{n,m}^*$), two *attributes* of website m from individual n 's perspective, namely, (*i*) whether website m reports a signal corresponding to the individual's prior ($y_m = x_n$), which gives desirability γ , and (*ii*) whether website m is identified as carrying a website majority signal ($y_m = z_n$), which gives desirability $1 - \gamma$. At the same time, website m 's rank ($r_{n,m}$) re-scales its final desirability in a multiplicatively separable way, and reflects individuals' tendency to devote more attention to higher ranked websites. Importantly, the multiplicative form leads individuals to trade off attributes and hence content, when making non-degenerate stochastic choices ($0 < \gamma < 1$). When $\gamma = 0$ or $\gamma = 1$ only one of the attributes matters.

Overall, $\rho_{n,m}$ summarizes some intuitive properties, namely, individuals tend to choose websites that are (*i*) higher ranked, (*ii*) ex-ante more informative, (*iii*) reflect their own prior information; and (*iv*) they may also trade off content and ranking.

2.4 Search Engine and Ranking Algorithm

The *ranking algorithm* 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.¹⁴

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,

$$r_{t,m} = \nu r_{t-1,m} + (1 - \nu) \rho_{t-1,m}, \quad (5)$$

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 (4). Such ranking dynamic reflects the algorithm used by search engines to

¹³Fortunato *et al.* (2006) and Demange (2014a) use related models of individuals' choices over ranked items. In section 6.1 we study a more flexible version of (4) that also parametrizes the degree of individuals' attention bias.

¹⁴For the sake of tractability, the ranking is considered as a cardinal score attached by the search engine to each given website. As pointed out by Demange (2014b, p. 918) a ranking is meant to measure the "relative strength of n items, meaning that the values taken by the scores matter up to a multiplicative constant".

update their rankings according to how “popular” a webpage is (Dean, 2013; MOZ, 2013; Vaughn, 2014).¹⁵ Accordingly, we will sometimes refer to the ranking as *popularity ranking*. The weight that is put on each term depends on the parameter ν , where for convenience we write,

$$(\nu, 1 - \nu) = \left(\frac{\kappa}{\kappa + 1}, \frac{1}{\kappa + 1} \right), \quad (6)$$

and where in turn $\kappa \in \mathbb{N}$ is a *persistence* parameter of the ranking algorithm. The larger κ , the more persistent the search engine’s ranking is. In Section 4, when studying asymptotic behavior, we will let $\kappa \rightarrow \infty$ as $N \rightarrow \infty$. This ensures that the effect of any given individual’s search on the ranking becomes vanishingly small and allows the dynamic process to converge to its unique limit.

2.5 Search Environments

A search environment is an *ex ante* notion that fixes the ranking algorithm, information structure and characteristics of individuals and websites, before they receive their signals and before they perform their search. More formally, we define a *search environment* \mathcal{E} as a list of variables,

$$\mathcal{E} = \langle (p, q, \mu); (N, \gamma); M; \kappa \rangle, \quad (7)$$

where (p, q, μ) describes the information structure, (N, γ) describes the individuals, M describes the websites, and κ describes the ranking algorithm. Given a search environment \mathcal{E} , we refer to an (*interim*) *realization* of \mathcal{E} as to the tuple $\langle \omega; (L, (y_m)_{m \in M}); ((x_n)_{n \in N}; (z_n)_{n \in N}) \rangle$, 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 $(x_n; z_n)$ 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 $m \in M$. For $J \subset M$, let $r_{t,J} = \sum_{m \in J} r_{t,m}$ and $\rho_{t,J} = \sum_{m \in J} \rho_{t,m}$ denote respectively total ranking and total clicking probability at time t on all websites in J ; we will be particularly interested in the case where $J = L$. 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 signals of agent n that enters to perform a search at time $t \equiv n$. (These expected probabilities are discussed in more detail in Appendix A.)

¹⁵This updating algorithm via the “popularity” of a website may be 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 vice versa).

3 Popularity Ranking and the Advantage of the Fewer (*AOF*)

The following proposition states our first main result. It illustrates a rather general phenomenon induced by popularity ranking, which we refer to as the *advantage of the fewer* (*AOF*), 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).¹⁶

Proposition 1. *Fix a search environment \mathcal{E} with uniform initial ranking r_1 , and N, κ large, and consider interim realizations of \mathcal{E} that vary in the number of outlets with correct signal. Then, the total limit clicking probability $(\rho_{N,J})$ on all outlets with a given signal (say $J \subset M$ where $m, m' \in J$ implies $y_m = y_{m'}$ and similarly for $m'', m''' \in M \setminus J$) is decreasing in the number of those outlets $(\#J)$, for “interior values” of the parameters (i.e., when $0 < \gamma < 1$, and at $J \neq 0, \frac{M-1}{2}, M-1$).*

This suggests that having *fewer* websites reporting a given signal *enhances* their overall traffic. The intuition behind the *AOF* effect is straightforward. When there are fewer websites with a given signal, the flow of individuals interested in reading about that signal are concentrated on fewer outlets, thus leading to relatively more clicks per outlet. This (trivial) static effect transforms into a dynamic, amplified one through the interaction of the popularity ranking with individuals’ stochastic choices: popularity ranking induces relatively higher rankings for those fewer websites, which makes them more attractive for subsequent individuals, inducing more trade-offs in favor of those fewer higher ranked websites, leading to even more clicks, and so on. The process, which can be seen as embodying a “*few get richer*” dynamic, gets repeated until it stabilizes in the limit, generating a potentially sizable amplification of total traffic on all websites with the given signal.¹⁷

The *AOF* implies that having a small majority (or a small minority) of websites rather than a large majority (or minority) of websites reporting a given signal actually increases their overall traffic. At the same time, the amount of traffic that such websites can attract is limited by the fact that majority (minority) websites attract all traffic from individuals with signals $x_n = z_n = y_K$ ($x_n = z_n = y_{M \setminus K}$). Accordingly, the results do not imply that all traffic is directed toward a single website or group of websites, but rather that there is an *amplification effect* for the traffic going to both majority or minority websites triggered by a lower number of corresponding websites. The *amplification effect* generated by the *AOF* hinges upon two main components: the popularity

¹⁶Throughout the paper, we use the term decreasing (and increasing) in the weak sense, that is, we say a function f is decreasing (increasing) if $x \geq y$ implies $f(x) \leq f(y)$ ($f(x) \geq f(y)$). When $x \in \mathbb{N}$, we say f is decreasing (increasing) at x if $f(x+1) \leq f(x)$ ($f(x+1) \geq f(x)$).

¹⁷Figure A.1 in the Online Appendix, provides a graphical illustration of the *amplification effect*.

ranking and a non-degenerate stochastic choice of the individuals, who trade off content depending on the ranking when γ is interior ($0 < \gamma < 1$). Instead, when $\gamma = 0$ or $\gamma = 1$ such an effect is absent and total traffic directed towards websites with a given signal is constant and does not depend on the number of websites (except if $J = 0, \frac{M-1}{2}, M - 1$).

Two further remarks regarding the *AOF* effect are in order.

1. *AOF* is potentially relevant for understanding the spread of misinformation, in the sense that a signal that is carried by few outlets because it is controversial or “fake news” may paradoxically receive amplified traffic *because* it is carried by few outlets.¹⁸ This is consistent with various claims that the algorithms used by Google and Facebook have apparently promoted websites reporting “fake news”.¹⁹ 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’ opinions and behavior (Carvalho *et al.*, 2011; Kata, 2012; Mocanu *et al.*, 2015; Shao *et al.*, 2016).
2. *AOF* holds beyond the specific assumptions of our model. The online appendix provides examples illustrating cases, where: (i) parameters are outside the assumed range, for example, when individual signals are uninformative ($p = \mu = 1/2$), (ii) rankings are not stochastic but rather a list of 1st to M^{th} ranked, as a function of traffic obtained ($\sum_{n' \leq n} \rho_{n'}$), (iii) individual choices are pure realizations of the stochastic choice functions (ρ_n), or (iv) where the total number of websites is not necessarily fixed, thus allowing for cases, where two or more websites can merge and thereby reduce the number of overall websites (M).²⁰ These examples further suggest that *AOF* is a potentially robust phenomenon.

¹⁸For example, if there was a single minority “fake news” website (so that $M - L = 1$), then the probability that that outlet would be visited by an individual in the limit is (assuming $\mu = 1$ and γ not too small) $\rho_{\infty, fake} = 1 - \frac{\gamma p(M-1)}{\gamma M - 1}$. Moreover, such a “fake news” website will be the top ranked one if $\gamma > \frac{1}{M(1-p)}$.

¹⁹On 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: “Three days after removing human editors, Facebook is already trending fake news,” *The Washington Post*, August 29, 2016.

²⁰One might expect that allowing for free entry would tend to weaken the *AOF*. However, while addressing free entry of websites and even modeling the strategic choices of websites is outside the scope of this paper, it is worth noting that the effect of allowing free entry on the *AOF* 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 websites (which are likely to also care about their reputation) to report such information. Also, new minority websites may not necessarily be able to “steal” much traffic from the existing ones due to their lower ranking (as explained by the *rich-get-richer* dynamic, see Section 6.2). Hence, we believe the result may be particularly relevant for dubious information, carried by relatively few websites and that “resonates” with a significant fraction of individuals.

4 Popularity Ranking and Asymptotic Learning

To assess the effect of the popularity ranking on opinion dynamics, we consider search environments from an *interim* and an *ex-ante* perspective. Since our model endogenizes individual clicking probabilities, it seems natural to evaluate efficiency in terms of asymptotic probability of clicking on a website carrying the correct signal. We further interpret this notion of efficiency as asymptotic learning, under the implicit assumption that individuals assimilate the content they read.²¹

Consider the probability of individual n choosing a website reporting a signal corresponding to the true state of the world ($y_m = \omega$). At the interim stage, we can write this as the probability $\rho_{n,L}$ ($= \sum_{m \in L} \rho_{n,m}$) of individual n clicking on any website $m \in L$. We can also define a measure of *interim efficiency* (\mathcal{P}_L), conditional on interim realizations, where the total number of websites reporting the correct signal is L , as:

$$\mathcal{P}_L(\alpha, \gamma, \mu, p) = \rho_{\infty,L} = \lim_{N \rightarrow \infty} \rho_{N,L}. \quad (8)$$

This implies the following measure of *ex ante efficiency* (\mathcal{P}):

$$\mathcal{P}(\alpha, \gamma, \mu, p, q) = \sum_{L=0}^M \binom{M}{L} q^L (1-q)^{M-L} \mathcal{P}_L(\alpha, \gamma, \mu, p), \quad (9)$$

which uses the accuracy of websites' signals (q) to weigh the different interim levels (\mathcal{P}_L).

To better highlight the role of popularity ranking on asymptotic learning, we use as benchmark of comparison, the case where ranking is random and uniform throughout ($r_{n,m} = \frac{1}{M}$ for all n, m), which we refer to simply as *random ranking*.

As explained in Appendix A, we compute the limit ranking and limit clicking probabilities using the mean dynamics approximation (Norman, 1972; Izquierdo and Izquierdo, 2013). This involves fixing an interim search environment (essentially characterized by the number of websites L carrying the correct signal) and approximating the random clicking probabilities $\rho_{n,m}$ in Equation (4) for $m \in M$ by their expectations $\hat{\rho}_{n,m} = \mathbb{E}[\rho_{n,m}]$, while letting $\kappa \rightarrow \infty$ as $N \rightarrow \infty$. 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 (9).

²¹For literature on asymptotic learning in models of social herding and social learning see Banerjee (1992); Bikhchandani *et al.* (1992, 1998); Smith and Sørensen (2000); Acemoglu *et al.* (2011); Acemoglu and Ozdaglar (2011).




Figure 1: *Interim* efficiency (\mathcal{P}_L) as a function of the number of websites with correct signal (L) for $\gamma = 0.33$ (black line) and other values of γ (dashed lines) for the cases of popularity ranking (left) and random ranking (right). In both panels, $M = 20$, $p = 0.55$, $\mu = 0.9$, and r_1 is uniform.

4.1 Non-monotonicity of Interim Efficiency

Before analyzing ex-ante efficiency, we first study interim efficiency. The following result follows from Proposition 1 and illustrates how interim efficiency is a non-monotonic function of the number of websites carrying the correct signal (L) due to the *advantage of the fewer* effect (*AOF*).

Corollary 1. *Fix a search environment \mathcal{E} with uniform initial ranking r_1 , and consider interim realizations of \mathcal{E} that vary only in the number of outlets with correct signal (L). Then \mathcal{P}_L is non-monotonic in L . In particular, when γ is interior ($0 < \gamma < 1$), then \mathcal{P}_L is increasing in L at $L = 0, \frac{M-1}{2}, M-1$, but it is decreasing in L otherwise.*

The non-monotonicity in L follows from three basic facts: (i) small majorities (or minorities) of outlets with a correct signal result in higher interim efficiency (\mathcal{P}_L) than large majorities (or minorities), resulting in a decreasing effect of L on \mathcal{P}_L induced by *AOF*, (ii) when L increases from $\frac{M-1}{2}$ to $\frac{M+1}{2}$, then interim efficiency increases, since the correct signal (y_L) switches from being a minority to being a majority signal, (iii) when L increases from 0 to 1 or from $M-1$ to M , interim efficiency obviously increases.

Figure 1 shows the non-monotonicity of interim efficiency \mathcal{P}_L in L , as stated in Corollary 1, as well as the *AOF* effect as discussed in Section 3, in the presence of the popularity ranking when the preference for like-minded news is interior ($\gamma = 0.33$), (black line, left panel), and shows that interim efficiency is monotonic and *AOF* is absent when $\gamma = 0, 1$ (left panel, dashed lines) or when the ranking is random (and therefore popularity ranking is switched off), (right panel).

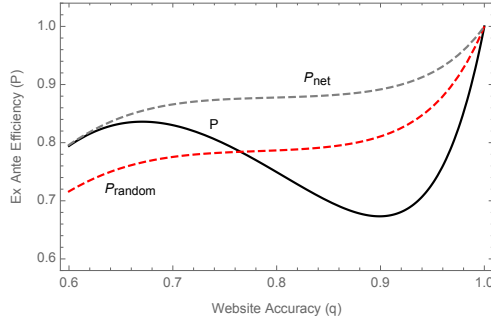


Figure 2: *Ex ante* efficiency (\mathcal{P}) as a function of the accuracy of websites' signals (q) (right) with popularity-ranking (black line) and random-ranking (red dashed). The figure also shows *ex ante* efficiency (gray dashed) with popularity-ranking. $M = 20$, $p = 0.55$, $\mu = 0.9$, $\gamma = 0.33$ and r_1 is uniform.

4.2 Ex Ante Efficiency

4.2.1 Comparative statics

Our notion of *ex ante* efficiency (\mathcal{P}) is a measure of asymptotic learning that obtains in our popularity-ranking based search environments. The following proposition describes comparative statics of \mathcal{P} with respect to basic parameters of the model.

Proposition 2. *Let \mathcal{E} be a search environment with a uniform initial ranking r_1 . Then:*

1. (*Individual accuracy*) \mathcal{P} is increasing in p and in μ .
2. (*Website accuracy*) \mathcal{P} can be both increasing or decreasing in q .
3. (*Like-mindedness*) \mathcal{P} is decreasing in γ , for $p \in (\frac{1}{2}, \bar{p}]$, for some $\bar{p} > \frac{1}{2}$.

We briefly discuss these effects.

1. (*Individual accuracy*) Higher levels of p and μ always increase interim hence also *ex ante* efficiency. This is not just the consequence of the direct effect of more accurate private signals. The direct effect is increased by the dynamic one, since more individuals receiving a correct signal, increases the number of clicks on websites reporting that signal, which in turn increases their ranking. This further increases the probability that subsequent individuals will also choose websites reporting a correct signal so that, even though individuals are assumed to be naïve, the higher p and μ are, the more effective popularity ranking is in aggregating information and creating a positive externality that enhances asymptotic learning.
2. (*Website accuracy*) A higher q may increase or decrease *ex ante* efficiency. This is a consequence of the non-monotonicity of the interim efficiency (Corollary 1), driven by *AOF*.

Higher values of q make it more likely that the number of websites with correct signal is large. Hence, due to *AOF*, a higher q decreases the clicking probability on websites with a correct signal, thus reducing ex ante efficiency for intermediate values of γ . If one could switch off *AOF*, then a higher q would always increase ex-ante efficiency (i.e., efficiency “net of *AOF*” is always increasing in q).²² This stark difference is illustrated in Figure 2.

3. (*Like-mindedness*) Given the assumption on the relative informativeness of private signals ($\mu q > p$), if moreover, p is not too large, then a higher γ decreases the probability of choosing a website reporting a correct signal.

4.2.2 Popularity-Ranking Effect (*PoR*)

We turn to the crucial question of how well popularity ranking performs in terms of asymptotic learning. The following definition uses random ranking as a comparison benchmark.

Definition 1. Let \mathcal{E} be a search environment with a uniform initial ranking r_1 , and let \mathcal{E}' be another search environment that differs from \mathcal{E} only in that the ranking is always an uniform random ranking. Let \mathcal{P} and \mathcal{P}' denote ex ante efficiency of \mathcal{E} and \mathcal{E}' respectively, then we define the popularity ranking effect of \mathcal{E} ($PoR(\mathcal{E})$) as the difference:

$$PoR(\mathcal{E}) = \mathcal{P} - \mathcal{P}'.$$

The following result compares popularity ranking and random ranking.

Proposition 3. Let \mathcal{E} be a search environment with uniform initial ranking r_1 . Then there exist $\bar{\gamma} > 0$, $\bar{\mu} < 1$, and a threshold function for q , $\phi : [0, 1] \rightarrow [\frac{1}{2}, 1]$, that is decreasing in $\hat{\gamma}$, such that $\phi(\hat{\gamma}) = 1$, for $\hat{\gamma} \in [0, \bar{\gamma}]$, and, moreover, for any threshold $\hat{\gamma} \in [\bar{\gamma}, 1]$:

1. $PoR(\mathcal{E}) \geq 0$ for $\gamma \in [0, \hat{\gamma}]$, $\mu \in [\bar{\mu}, 1]$, provided $q \in [\frac{p}{\mu}, \phi(\hat{\gamma})]$.
2. $PoR(\mathcal{E}) \leq 0$ for $\gamma \in [\bar{\gamma}, \hat{\gamma}]$, provided $q \in [\phi(\hat{\gamma}), 1]$.

²²Specifically, one can define ex ante efficiency “net of *AOF*” (\mathcal{P}_{net}) as ex ante efficiency calculated with weighted constant “average” minority and majority traffic levels, respectively, $\mathcal{P}_{\lceil \frac{M}{4} \rceil}$ and $\mathcal{P}_{\lceil \frac{3M}{4} \rceil}$, formally:

$$\mathcal{P}_{net} = \sum_{k=1}^{\frac{M-1}{2}} \mathcal{P}_{\lceil \frac{M}{4} \rceil} \binom{M}{k} q^k (1-q)^{M-k} + \sum_{k=\frac{M+1}{2}}^{M-1} \mathcal{P}_{\lceil \frac{3M}{4} \rceil} \binom{M}{k} q^k (1-q)^{M-k} + q^M.$$

While, in principle, it would be possible to correct the ranking algorithm for the *AOF* effect, we are unaware of any such correction undertaken in practice, nor have we seen the effect mentioned in the computer science or machine learning literature.

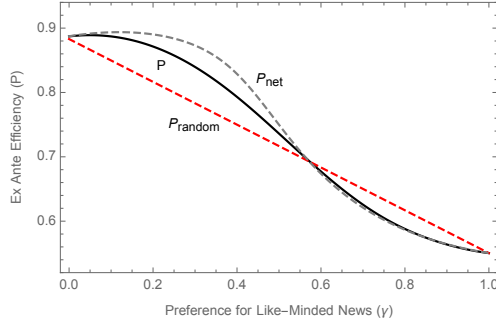


Figure 3: *Ex ante* efficiency (\mathcal{P}) as a function of the preference for like-minded news (γ) with popularity-ranking (black line) and random-ranking (red dashed). The figure also shows net *ex ante* efficiency (gray dashed) with popularity-ranking. The plot is drawn for $M = 20$, $p = 0.55$, $\mu = 0.9$, $q = 0.7$, and r_1 is uniform.

In other words, when preference for like minded news is sufficiently low ($\gamma \in [0, \bar{\gamma}]$) and the accuracy on the majority signal is sufficiently high ($\mu \in [\bar{\mu}, 1]$), then popularity ranking does better than random ranking. That is, when individual clicking behavior generates sufficiently positive information externalities, popularity ranking aggregates information and increases the probability of asymptotic learning relative to random ranking. Moreover, this continues to hold for higher values of preference for like minded news ($\gamma \in [\bar{\gamma}, \hat{\gamma}]$, for $\hat{\gamma}$ that can go up to 1), provided the accuracy of websites is not too high ($q \leq \phi(\hat{\gamma})$) (due to *AOF*). Put differently, due to *AOF*, an increase in the ex-ante informativeness of websites (q), may lead to a negative popularity ranking effect.²³ Figures 2 and 3 illustrate Proposition 3: popularity ranking dominates random ranking in terms of ex-ante efficiency provided q and γ are not too large, and is dominated by random ranking otherwise.

5 Personalized Ranking

Another important question for a ranking algorithm concerns whether it should keep track of, and use, information it has available concerning the individuals’ identity and past searches. A search engine may want its algorithm to condition the outcomes of searches on the geographic location of the individuals (e.g., using the individual’s *IP address*) or other individual characteristics (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.²⁴

²³By contrast, ex ante efficiency “net of *AOF*” (\mathcal{P}_{net}) is above random ranking even for large values of q .

²⁴See Pariser (2011); Dean (2013); MOZ (2013); Vaughn (2014); Kliman-Silver *et al.* (2015). Hannak *et al.* (2013) document the presence of extensive personalization of search results. In particular, while they show that the extent of 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 personalization based on the *Booble* extension of Chrome.

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 $N_A = \#A > 0$ and $N_B = \#B > 0$. A *personalized ranking algorithm* 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 different than those from the other one. Set, for any t and m , and for $\ell = A, B$,

$$r_{t,m}^\ell = \nu_t^\ell r_{t-1,m}^\ell + (1 - \nu_t^\ell) \rho_{t-1,m}^\ell, \quad (10)$$

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:

$$(\nu_t^\ell, 1 - \nu_t^\ell) = \left(\frac{\kappa}{1 - \lambda_t + \kappa}, \frac{1 - \lambda_t}{1 - \lambda_t + \kappa} \right), \text{ where } \lambda_t = \begin{cases} 0 & \text{if } t - 1 \in \ell \\ \lambda & \text{else,} \end{cases} \quad (11)$$

with $\kappa \in \mathbb{N}$ and $\lambda \in [0, 1]$ parameters of the personalized ranking algorithm. 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 $\frac{\kappa}{1-\lambda+\kappa}, \frac{1-\lambda}{1-\lambda+\kappa}$). In particular, 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 algorithm partially separates individuals according to some individual characteristics (e.g., individuals in different groups may have different parameters γ). We will refer to a *personalized search environment* \mathcal{E}_λ when considering the generalization of a search environment \mathcal{E} , defined in Section 2.5, to the case where the ranking is described by Equations (10) and (11) with $\lambda \geq 0$.

5.1 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 $\lambda = 0$, there is no difference in the ranking of websites for the two groups, while as λ increases the groups start observing potentially different rankings which may further trigger different website choices, thus leading to different opinions. In other words, increased personalization may lead to increased *belief polarization*.

Definition 2. Fix a search environment \mathcal{E} with nonempty groups of individuals A and B . 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 more belief polarization than \mathcal{E}' , if $\mathcal{BP}(\mathcal{E}) > \mathcal{BP}(\mathcal{E}')$.

When individuals are separated according to their preference for like-minded news, a larger personalization parameter λ will lead to more polarization as measured by \mathcal{BP} .²⁵

Proposition 4. *Let \mathcal{E}_λ be a personalized search environment with personalization parameter λ . Suppose there are two groups of individuals A and B of equal size and that have different preference for like-minded news ($\gamma_A \neq \gamma_B$). Then, $\mathcal{BP}(\mathcal{E}_\lambda)$ is increasing in λ .*

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).

While belief polarization, as a feature of personalized ranking, may be deemed undesirable, in the next section we further analyze personalized ranking in terms of ex ante efficiency.

5.2 Personalized-Ranking Effect (*PeR*)

We turn to the effects of personalization of search results on ex-ante efficiency. The following proposition focuses on the case where individuals in different groups have different levels of preferences for like-minded news (say γ_A and γ_B). Personalization here can be seen as progressively “separating” the two groups by uncoupling their rankings and thereby switching off potential externalities from one group to the other. To the extent that the group with weaker preference for like minded news exerts a positive externality on the groups’ rankings and overall ex ante efficiency, increasing the personalization may inhibit overall ex ante efficiency.

Definition 3. *Let \mathcal{E} be a non-personalized and \mathcal{E}_λ be a personalized search environment with personalization parameter λ , and both with a uniform initial ranking r_1 . Let \mathcal{P} and \mathcal{P}_λ denote ex*

²⁵Notice 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.

ante efficiency of \mathcal{E} and \mathcal{E}_λ respectively, then we can define the personalized ranking effect of \mathcal{E}_λ ($PeR(\mathcal{E}_\lambda)$) as the difference:

$$PeR(\mathcal{E}_\lambda) = \mathcal{P}_\lambda - \mathcal{P}.$$

Proposition 5. \mathcal{E}_λ be a personalized search environment with personalization parameter λ , and with a uniform initial ranking r_1 . Suppose there are two nonempty groups of individuals A and B differentiated with respect to γ , with $0 \leq \gamma_A < \gamma_B \leq \bar{\gamma}$. Then there exist $\bar{\gamma} > 0$, $\bar{\mu} < 1$, and a decreasing function of $\hat{\gamma}$, $\phi : [0, 1] \rightarrow [\frac{1}{2}, 1]$, such that $\phi(\hat{\gamma}) = 1$, and, moreover, for any $\hat{\gamma} \in [\bar{\gamma}, 1]$:

1. $PeR(\mathcal{E}) \leq 0$ for $\gamma \in [0, \hat{\gamma}]$, $\mu \in [\bar{\mu}, 1]$, provided $q \in [\frac{p}{\bar{\mu}}, \phi(\hat{\gamma})]$.
2. $PeR(\mathcal{E}) \geq 0$ for $\gamma \in [\bar{\gamma}, \hat{\gamma}]$, provided $q \in [\phi(\hat{\gamma}), 1]$.

Although the exact cutoff values for the parameters need not coincide, the parallels between the popularity ranking effect and the personalized ranking effect are stark. That is, when the preference for like minded news is sufficiently low ($\gamma \in [0, \bar{\gamma}]$) and the accuracy on the majority signal is sufficiently high ($\mu \in [\bar{\mu}, 1]$), then personalized ranking does worse than non-personalized (popularity) ranking; moreover, this continues to hold for higher values of preference for like minded news ($\hat{\gamma} \in [\bar{\gamma}, 1]$), provided the accuracy of websites is not too high ($q \leq \phi(\hat{\gamma})$) (due to *AOF*); on the other hand for the higher values of preference for like-minded news ($\hat{\gamma} \in [\bar{\gamma}, 1]$) personalized ranking will perform better than non-personalized (popularity) ranking, provided accuracy of websites is sufficiently high ($q \geq \phi(\hat{\gamma})$). In other words, the forces that lead to a positive (negative) popularity effect are similar to the ones that lead to a negative (positive) personalization effect.

The intuition for the negative relation between the popularity ranking and the personalized ranking effects (*PoR* and *PeR*) is due to the fact that a positive popularity ranking effect occurs when individuals' parameters generate positive feedback into the dynamics (high accuracy μ and low preference for like-minded news γ); the same forces favor non-personalization and hence tend to generate a negative personalized ranking effect. This is because personalization can be seen as limiting the feedback between individuals in the dynamics, and so, when individuals' signals tend to generate positive feedback, it is better not to limit the feedback and hence it is better not to personalize the ranking and *vice versa* when individuals' signals generate negative feedback. The similarity also with respect to the website accuracy parameter (q) is due to the *AOF* effect that is present with or without personalization.

Figure 4 illustrates the negative relationship between *PoR* and *PeR*. In this case, personalized ranking (green dashed) tends to be dominated (in terms of asymptotic learning) by either non-personalized ranking (black) or by random ranking (gray dashed). The right-hand panel indicates




Figure 4: *Ex ante* efficiency with no personalization ($\mathcal{P}_{\lambda=0}$) (black) and with personalization w.r.t. γ ($\mathcal{P}_{\lambda=1}$) (green dashed); and with random ranking (gray dashed). The right panel illustrates the same for ex-ante efficiency “net of *AOF*”. In both panels, $M = 20$, $q = 0.7$, $p = 0.55$, $\mu = 0.9$, $\gamma = 0.33$, and r_1 uniform.

that when *AOF* is switched off, the corresponding “net” *PeR* effect tends to be negative for any level of ex-ante accuracy of websites.

It is important to note that our searches are common value searches in the sense that all individuals want to read outlets carrying the (same) correct signal. When individuals have private values, personalization may be an important tool that actually favors asymptotic learning of typically multiple and distinct signals.

6 Robustness

6.1 Attention bias

In our basic model of how individuals respond to website rankings, we assumed multiplicative separability between the ranking (r_n) and the ranking-free values (v_n^*), such that the ranking entered directly as a weighting function for the ranking-free-values. We now generalize the weighting function by introducing an *attention bias* parameter $\alpha \geq 0$, which yields the following more general ranking-weighted values:

$$v_{n,m} = (r_{n,m})^\alpha \cdot v_{n,m}^*.$$

The parameter α calibrates the individual’s attention bias in the following sense: $\alpha = 1$ is a neutral benchmark in that it maintains the weight differences already present in the entries of the ranking r_n ; $\alpha > 1$ magnifies the differences in the entries of r_n , and in particular also the ones present in the initial ranking r_1 ; $\alpha < 1$ reduces the differences in the entries of r_n ; in the limit, as $\alpha \rightarrow 0$, all entries have the same weight, which represents the case with *no* attention bias, where all websites that provide the same signal yield the same value. Given the values $v_{n,m}$, we can write the *website*

choice function with attention bias $\alpha \geq 0$, $\rho_n : \{0, 1\}^2 \times \Delta(M) \rightarrow \Delta(M)$, as

$$\rho_{n,m} = \frac{(r_{n,m})^\alpha \cdot v_{n,m}^*}{\sum_{m' \in M} (r_{n,m'})^\alpha \cdot v_{n,m'}^*}. \quad (12)$$

When $\alpha = 1$, this coincides with the choice function studied so far (Equation (4)). When $\alpha = 0$, it yields choices that coincide with the ranking-free choices $v_{n,m}^*$, since all outlets have the same weight $(r_{n,m})^0 = 1$. For this same reason, it is easy to see that, in terms of website choices and hence interim and ex ante efficiency (\mathcal{P}_L and \mathcal{P}), an environment with popularity ranking $r_{n,m}$ but attention bias $\alpha = 0$ is outcome-equivalent to an environment with *random ranking*, regardless of what α might be in that environment.²⁶

The *AOF* effect stated in Proposition 1 (i.e., for the case of $\alpha = 1$) carries over *verbatim* to the general case of $\alpha \geq 0$,²⁷ and similarly its implications for asymptotic learning or interim efficiency (Corollary 1), as summarized by the following Corollary:

Corollary 2. *Fix a search environment \mathcal{E} with attention bias parameter $\alpha \geq 0$ and with uniform initial ranking r_1 and κ large, and consider interim realizations of \mathcal{E} that vary only in the number of outlets with correct signal (L). Then:*

1. *For $\alpha > 0$, \mathcal{P}_L is non-monotonic in L .*
2. *For $\alpha = 0$, \mathcal{P}_L is monotonically increasing in L .*

In particular, when $\alpha > 0$ and $0 < \gamma < 1$, then \mathcal{P}_L is increasing in L at $L = 0, L = M - 1$, as well as at $L = \frac{M-1}{2}$, but it is decreasing in L otherwise.

In particular, this shows that *AOF* applies when ranking matters ($\alpha > 0$) and disappears when there is no attention bias and ranking does not matter ($\alpha = 0$). This also implies that *AOF* disappears with random ranking. In terms of comparative statics of ex ante efficiency, the results of Proposition 2 also carry over *verbatim* to the general case of $\alpha \geq 0$, with the only difference that, when $\alpha = 0$, then \mathcal{P} is increasing in q always, while it can be increasing or decreasing in q for $\alpha > 0$; again this is because *AOF* applies when $\alpha > 0$, but disappears when $\alpha = 0$.

²⁶When there is zero attention bias ($\alpha = 0$), all websites receive equal weights in the website values ($v_{n,m}$) that determine the website choice probabilities ($\rho_{n,m}$). Similarly, when there is random ranking (r_1 and every subsequent ranking r_n is uniform, so that $r_{n,m} = \frac{1}{M}$ regardless of the parameter α), then all websites will also receive equal weights in the website values. Thus website choice probabilities ($\rho_{n,m}$) always coincide in the two cases.

²⁷The proof of Proposition 1 in Appendix B is directly given for the case of $\alpha \geq 0$. It is important to note that, while, for $0 \leq \alpha \leq 1$, the unique limit always satisfies *AOF*; for $\alpha > 1$, depending on the initial condition (i.e., the initial ranking; see also Section 6.2), there can be multiple limits, of which the asymptotically stable ones, that is, the only ones that can be reached by our dynamic process, also always satisfy *AOF*.

6.2 Non-Uniform Initial Ranking and the Rich Get Richer

Many of the propositions stated in the paper assumed a uniform initial ranking. We now study this assumption in more detail, looking at environments with general attention bias ($\alpha \geq 0$). As it turns out, the expected limit clicking probabilities $\hat{\rho}_{\infty,m}$ only depend on the initial ranking when $\alpha > 1$. In this case, there is also a “rich-get-richer” effect.

Proposition 6. *Let \mathcal{E} be a search environment with attention bias parameter $\alpha \geq 0$ and with interim realization $\langle \omega; (L, (y_m)) \rangle$, with interior initial ranking $r_{1,m} > 0$ for all m . Then, if attention bias satisfies $0 \leq \alpha \leq 1$, then the expected limit clicking probabilities $\hat{\rho}_{\infty,m}$ do not depend on r_1 . This is not true if $\alpha > 1$.*

The evolution of a website’s ranking based on its “popularity,” interacted with a sufficiently large attention bias ($\alpha > 1$) 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 the ranking probabilities (r_n) and in the expected website choice probabilities ($\hat{\rho}_n$), are driven by the initial ranking (r_1) and are amplified by the attention bias. We define this more formally.

Definition 4. *Fix a search environment \mathcal{E} with attention bias parameter $\alpha \geq 0$ and 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 7. *Let \mathcal{E} be a search environment with attention bias parameter $\alpha \geq 0$ 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 $n > 1$:*

$$\frac{\hat{\rho}_{n,m}}{\hat{\rho}_{n,m'}} \begin{cases} < \\ = \\ > \end{cases} \frac{\hat{\rho}_{n-1,m}}{\hat{\rho}_{n-1,m'}} \begin{cases} \text{if } \alpha < 1 \\ \text{if } \alpha = 1 \\ \text{if } \alpha > 1. \end{cases} \quad (13)$$

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

This proposition shows that the attention bias 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 websites with the same signal will tend to be visited with the same probability in the limit. When $\alpha = 1$ 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. 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).²⁸

6.3 Sophisticated Learning

The main contribution of our model is to combine endogenous ranking of websites with sequential clicking behavior of behaviorally biased individuals. So far, we assumed individuals that are naïve with regards to the ranking. They do not make inferences on the informativeness of websites’ signals based on the observed rankings. Accordingly, our theoretical framework does not encompass a fully-fledged model of rational herding (Banerjee, 1992; Bikhchandani *et al.*, 1992; Smith and Sørensen, 2000). However, sophisticated individuals sufficiently far ahead in the sequence (i.e., n large) might exploit the website ranking to infer information on the private signals of other individuals. The following three remarks discuss some implications of this for our analysis.

1. Our model implicitly embeds a reduced-form version of rational herding. Individuals have a tendency to choose higher ranked websites. Consequently, as seen in Section 4, depending on the underlying parameters of the environment, the interaction between individual stochastic choices and the popularity ranking can induce information aggregation in the limit. The more people receive a correct signal, the more likely it is that websites with a correct signal will be chosen, the higher their ranking and the more likely it is that subsequent individuals will also choose websites reporting such a signal, and so on.

²⁸Indeed, 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.




Figure 5: *Interim* efficiency (\mathcal{P}_L) with “rationally computed signals” (\hat{z}_n) as a function of the number of websites with correct signal (L) for $\gamma = 0.33$ (black line) and other values of γ (dashed lines) for the cases of approximate signals ($\hat{\mu} = 0.9$) (left) and fully accurate signals ($\hat{\mu} = 1$) (right). In both panels, $M = 20$, $p = 0.55$ and r_1 is uniform.

2. A general model of rational herding would require a rather sophisticated belief updating process from the perspective of individuals. In particular, individual t would need to calculate the probability of individuals $t - 1, t - 2, \dots$, having received a private signal equal to the website-majority signal, given: (i) the observed ranking at time t , (ii) the prior beliefs on the initial ranking at time $t = 1$, and (iii) the beliefs about t itself (i.e., the individual’s own position in the search sequence). 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).²⁹
3. To get a better idea of the limiting behavior that would obtain with more sophisticated individuals, suppose that with some degree of approximation individuals can keep track of the items mentioned above. Then, given the parameters of the model (and given $\mu q > p > \frac{1}{2}$), for sufficiently large n , individuals can compute which outlets have been selected by a *majority* of individuals following their signal x_t (for $t < n$). For sufficiently large n and sufficiently small computation costs, the “rational side” of individuals’ choices (with desirability value $1 - \gamma$) will be to ignore both their own signals x_n and z_n and to choose their computed signal

²⁹Several 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-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.

of what a majority of individuals following their signal x_t have chosen so far ($t < n$). Call this computed signal \hat{z}_n , and suppose that, for sufficiently large n , computation and observation costs are such that this can only be observed or deduced with an accuracy $\hat{\mu} > \mu q > p$. Then, the signal \hat{z}_n is eventually used instead of the majority signal z_n . Figure 5 illustrates the limiting behavior when the signal \hat{z}_n is approximate ($\hat{\mu} = 0.9$) (left) and when it is fully accurate ($\hat{\mu} = 1$) (right). As can be seen, in both cases, when the preference for like-minded news is interior ($\gamma = 0.33$, black line), the *AOF* effect continues to hold, but *without* the upward jump when L goes from $\frac{M-1}{2}$ to $\frac{M+1}{2}$, since the website majority signal, responsible for such a jump, no longer plays a role when the signal \hat{z}_n is sufficiently accurate. Comparing the two figures to Figure 1 (left panel), it is clear that the largest discrepancy in terms of interim efficiency occurs when the website majority-signal is not correct (i.e., outlets in L are a minority), in which case the naïve individual’s “rational” choice of following the ex-ante most informative signal (z_n) is clearly sub-optimal.³⁰

7 Conclusions

The economics literature has largely overlooked the informational role played by ranking algorithms used by platforms such as search engines or social media to rank and filter the virtually infinite amount of information potentially available to their users. To address this neglected issue, we developed a model that allows us to study how basic features of a search engine’s ranking algorithm—popularity and personalization—interact with individual search behavior and how this affects website traffic, opinion dynamics and asymptotic learning.

Our results uncovered a rather general property of popularity-based rankings, we refer to as the *advantage of the fewer (AOF)*. Roughly speaking, it suggests that, given individuals’ attention bias and preferences over content, the smaller the number of websites reporting a given information, the larger the share of traffic directed to those fewer websites. Because dubious or particularly controversial information is often carried by a small number of websites, the *AOF* effect may help explain the spread of misinformation, since it shows how being small in number may actually boost overall traffic to such websites. Nevertheless, despite the *AOF* effect, popularity rankings can have an overall positive effect on asymptotic learning by fostering information aggregation, as long as individuals can, on average, provide sufficiently positive feedback to the ranking algorithm. The

³⁰Notice also that a ranking algorithm that keeps track of all clicks can potentially learn a lot about the environment and the true state of the world. 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.

model further provides insights on a controversial component of the ranking algorithm, namely personalization. While personalized rankings can clearly be efficient for search queries on private value issues (e.g., where to have dinner), we find that for queries on common value issues (e.g., whether or not to vaccinate a child) they can hinder asymptotic learning and can also deepen *belief polarization*.

So far, public authorities have worried about the possible strategic manipulation of website rankings by search engines regarding transactional queries.³¹ Our results suggest that, even in a “neutral 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; Hagiú 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 asymptotic learning.

³¹http://europa.eu/rapid/press-release_MEMO-15-4781_en.htm

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.
- AGRANOV, M. and ORTOLEVA, P. (2017). Stochastic choice and preferences for randomization. *Journal of Political Economy*, **125** (1), 40–68.
- 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.
- BLOCK, H. D. and MARSCHAK, J. (1960). Random orderings and stochastic theories of responses. In I. Olkin, S. Ghurye, W. Hoeffding, W. Madow and H. Mann (eds.), *Contributions to probability and statistics*, vol. 2, Stanford University Press Stanford, CA, pp. 97–132.
- BOHREN, J. A. (2016). Informational with Model Misspecification. *Journal of Economic Theory*, **163**, 222–247.
- 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.
- 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. (2014a). Collective Attention and Ranking Methods. *Journal of Dynamics and Games*, **1** (1), 17–43.
- (2014b). A ranking method based on handicaps. *Theoretical Economics*, **9** (3), 915–942.
- 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. 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.
- 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.
- GÜL, F., NATENZON, P. and PESENDORFER, W. (2014). Random choice as behavioral optimization. *Econometrica*, **82** (5), 1873–1912.
- 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.
- LUCE, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.
- MENCZER, F., FORTUNATO, S., FLAMMINI, A. and VESPIGNANI, A. (2006). Googlearchy or Googleocracy. *IEEE Spectrum Online*.
- MOCANU, D., ROSSI, L., ZHANG, Q., KARSAI, M. and QUATTROCIOCCI, 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.
- 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.
- 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.
- 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,

$$\begin{aligned} r_{t,m} &= \nu_t r_{t-1,m} + (1 - \nu_t) \rho_{t-1,m} \\ &= \frac{\kappa_t}{1 + \kappa_t} r_{t-1,m} + \frac{1}{1 + \kappa_t} \rho_{t-1,m} \\ &= r_{t-1,m} + \frac{1}{1 + \kappa_t} (\rho_{t-1,m} - r_{t-1,m}) \end{aligned}$$

Then, given Equation (12), for search environments with attention bias (α),

$$r_{t,m} = r_{t-1,m} + \frac{1}{1 + \kappa_t} \left(\frac{(r_{t-1,m})^\alpha \cdot v_{t-1,m}^*}{\sum_{m'} (r_{t-1,m'})^\alpha \cdot v_{t-1,m'}^*} - r_{t-1,m} \right),$$

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 given by the expressions $v_{t-1,m}^*$. 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

$$\begin{aligned} \hat{\rho}_{t-1,m} = \mathbb{E}[\rho_{t-1,m}] &= p\mu \frac{(\hat{r}_{t-1,m})^\alpha \cdot \hat{v}_m^{*00}}{\sum_{m'} (\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*00}} + p(1 - \mu) \frac{(\hat{r}_{t-1,m})^\alpha \cdot \hat{v}_m^{*01}}{\sum_{m'} (\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*01}} \\ &+ (1 - p)\mu \frac{(\hat{r}_{t-1,m})^\alpha \cdot \hat{v}_m^{*10}}{\sum_{m'} (\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*10}} + (1 - p)(1 - \mu) \frac{(\hat{r}_{t-1,m})^\alpha \cdot \hat{v}_m^{*11}}{\sum_{m'} (\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*11}} \quad (\text{A.1}) \end{aligned}$$

is the expected clicking probability of the individual entering in period $t-1$, and where the expected ranking-free values $\hat{v}_m^{*00}, \hat{v}_m^{*01}, \hat{v}_m^{*10}, \hat{v}_m^{*11}$ are given by the following table:

	\hat{v}_m^{*00}	\hat{v}_m^{*01}	\hat{v}_m^{*10}	\hat{v}_m^{*11}
$m \in L, y_L \neq y_K$	γ/L	$1/L$	0	$(1 - \gamma)/L$
$m \in L, y_L = y_K$	$1/L$	γ/L	$(1 - \gamma)/L$	0
$m \notin L, y_L \neq y_K$	$(1 - \gamma)/(M - L)$	0	$1/(M - L)$	$\gamma/(M - L)$
$m \notin L, y_L = y_K$	0	$(1 - \gamma)/(M - L)$	$\gamma/(M - L)$	$1/(M - L)$

for $L \neq 0, M$.³²

Here \hat{v}_m^{*00} represents the expected probability of an individual choosing a website m , contingent on having received two correct signals $x_n = \omega$ and $z_n = y_K$, when $\alpha = 0$ (i.e., absent attention bias); \hat{v}_m^{*01} represents the expected probability of an individual choosing a website m , contingent on having received a correct signal on the state of the world, $x_n = \omega$, and an incorrect signal on the majority, $z_n \neq y_K$, when $\alpha = 0$, and analogously for \hat{v}_m^{*10} and \hat{v}_m^{*11} . Importantly, $\hat{v}_m^{*00}, \hat{v}_m^{*01}, \hat{v}_m^{*10}$, and \hat{v}_m^{*11} 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, $\frac{1}{1+\kappa_t} \rightarrow 0$ as $t \rightarrow \infty$, 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 $x = (x_1, \dots, x_M)$), we obtain a function $g : \Delta_\epsilon(M) \rightarrow \mathbb{R}^M$, defined, for $m = 1, 2, \dots, M$, 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,

$$\begin{aligned} \theta_m(x) = & p\mu \frac{(x_m)^\alpha \cdot \hat{v}_m^{*00}}{\sum_{m'} (x_{m'})^\alpha \cdot \hat{v}_{m'}^{*00}} + p(1-\mu) \frac{(x_m)^\alpha \cdot \hat{v}_m^{*01}}{\sum_{m'} (x_{m'})^\alpha \cdot \hat{v}_{m'}^{*01}} + \\ & + (1-p)\mu \frac{(x_m)^\alpha \cdot \hat{v}_m^{*10}}{\sum_{m'} (x_{m'})^\alpha \cdot \hat{v}_{m'}^{*10}} + (1-p)(1-\mu) \frac{(x_m)^\alpha \cdot \hat{v}_m^{*11}}{\sum_{m'} (x_{m'})^\alpha \cdot \hat{v}_{m'}^{*11}}. \end{aligned}$$

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 (and hence also ρ_t) tends to follow the unique solution trajectory and linger around the asymptotically stable limit point of the differential equation $\dot{x} = g(x)$. (See Izquierdo and Izquierdo (2013), results (i) and (iii) on p. 261).

B Proofs

Proof of Proposition 1. To avoid duplication, we prove the proposition directly for the case of $\alpha > 0$ mentioned in Section 6.1, which includes the case of $\alpha = 1$, stated in Proposition 1, as a special case.

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, $\hat{r}_{n,m} = \hat{r}_{n,m'}$ for any two websites m, m' with $y_m = y_{m'}$. Since we consider partitions of M into sets J and $M \setminus J$, and $\rho_{n,J} + \rho_{n,M \setminus J} = 1$, it suffices to check the case $J = L$. To simplify notation, let $x = \hat{r}_{n,L} = L \cdot \hat{r}_{n,m}$ for $m \in L$ be the total expected ranking probability for the websites in L , so that $1 - x$ is the total expected ranking probability for the remaining websites in $M \setminus L$. Using the expressions for $\hat{v}_m^{*00}, \hat{v}_m^{*01}, \hat{v}_m^{*10}, \hat{v}_m^{*11}$ from the table in Appendix A, and assuming that $0 \leq x \leq 1$, we have that the equations defining the

³²In the extreme cases, $L = 0$ or $L = M$, clearly, $\hat{v}_m^{*00} = \hat{v}_m^{*01} = \hat{v}_m^{*10} = \hat{v}_m^{*11} = 0$ ($= 1$) if $m \in L, L = 0$ or $m \notin L, L = M$ (if $m \in L, L = M$ or $m \notin L, L = 0$).

limit probabilities can be reduced to two equations of the form:

$$H_L^{\text{minority}}(x; \alpha, \mu, \gamma, p) \equiv \theta_L^{\text{minority}}(x; \alpha, \mu, \gamma, p) - x = 0, \text{ for } 1 \leq L \leq \frac{M-1}{2} \quad (\text{B.1})$$

$$H_L^{\text{majority}}(y; \alpha, \mu, \gamma, p) \equiv \theta_L^{\text{majority}}(y; \alpha, \mu, \gamma, p) - y = 0, \text{ for } \frac{M-1}{2} \leq L \leq M-1, \quad (\text{B.2})$$

where $\theta_L^{\text{minority}}$ and $\theta_L^{\text{majority}}$ are defined as:

$$\theta_L^{\text{minority}}(x; \alpha, \mu, \gamma, p) = p(1-\mu) + \frac{p\mu\gamma\left(\frac{x}{L}\right)^\alpha}{\gamma\left(\frac{x}{L}\right)^\alpha + (1-\gamma)\left(\frac{1-x}{M-L}\right)^\alpha} + \frac{(1-p)(1-\mu)(1-\gamma)\left(\frac{x}{L}\right)^\alpha}{(1-\gamma)\left(\frac{x}{L}\right)^\alpha + \gamma\left(\frac{1-x}{M-L}\right)^\alpha} \quad (\text{B.3})$$

and

$$\theta_L^{\text{majority}}(y; \alpha, \mu, \gamma, p) = p\mu + \frac{p(1-\mu)\gamma\left(\frac{y}{L}\right)^\alpha}{\gamma\left(\frac{y}{L}\right)^\alpha + (1-\gamma)\left(\frac{1-y}{M-L}\right)^\alpha} + \frac{(1-p)\mu(1-\gamma)\left(\frac{y}{L}\right)^\alpha}{(1-\gamma)\left(\frac{y}{L}\right)^\alpha + \gamma\left(\frac{1-y}{M-L}\right)^\alpha} \quad (\text{B.4})$$

It is easy to check that at $x = 0$ and $x = 1$, we have, respectively,

$$0 \leq p(1-\mu) = \theta_L^{\text{minority}}(0; \alpha, \mu, \gamma, p) < \theta_L^{\text{minority}}(1; \alpha, \mu, \gamma, p) = p + (1-p)(1-\mu) < 1,$$

and similarly, at $y = 0$ and $y = 1$,

$$0 < p\mu = \theta_L^{\text{majority}}(0; \alpha, \mu, \gamma, p) < \theta_L^{\text{majority}}(1; \alpha, \mu, \gamma, p) = p + (1-p)\mu \leq 1.$$

In particular, $\theta_L^{\text{minority}}$ starts at or above the x -function at $x = 0$ and ends below the x -function at $x = 1$. Similarly, $\theta_L^{\text{majority}}$ starts above the y -function at $y = 0$ and ends below the y -function at $y = 1$. In Figure B.1, we plot $\theta_L^{\text{majority}}$ for different values of L and for $\alpha = 1$ on the panel on the left and for $\alpha = 4$ on the panel on the right. For most parameter values of interest (i.e., while α not too large) there is a unique interior solution to both Equations (B.1) and (B.2); this is the situation depicted on the left panel. However, in general there can be multiple solutions; as depicted in the right panel. Importantly, the solutions of interest are the ones where the functions $\theta_L^{\text{minority}}$ and $\theta_L^{\text{majority}}$ intersect the x - and y -functions from above.³³ This is because, the process can only converge to those due to the fact that while clicking probabilities ($\theta_L^{\text{minority}}$ or $\theta_L^{\text{majority}}$) are above the respective ranking probabilities (x or y) then the ranking probabilities will tend to increase and this will continue until the solution (at the intersection) is reached. Similarly, when clicking probabilities are below the respective ranking probabilities, then the ranking probabilities will tend to decrease until the solution is reached. In particular, only the two solutions with arrows are relevant in the panel on the right. The other interior solutions without an arrow (where the x - and y -functions are intersected from below) are unstable and are never reached by our dynamic process.³⁴ The key part of the proof is to show that the (stable) solutions $x = x_L^{\text{minority}}$ and

³³As $\alpha \rightarrow \infty$, the stable solutions converge to the bounds $p(1-\mu)$ and $p + (1-p)(1-\mu)$ for $\theta_L^{\text{minority}}$ and to $p(1-\mu)$ and $p + (1-p)\mu$ for $\theta_L^{\text{majority}}$, where both functions, $\theta_L^{\text{minority}}$ and $\theta_L^{\text{majority}}$, become arbitrarily flat. Because these solutions do not depend on L , this implies that the AOF effect tends to zero in the limit. On the other hand, the unstable solutions converge to $\frac{L}{M}$, where both $\theta_L^{\text{minority}}$ and $\theta_L^{\text{majority}}$ become arbitrarily steep.

³⁴Even if the process were to start at such a solution, any extra click by an individual will move the ranking probability to the left (or to the right) thereby leading to a situation where the clicking probability is below the ranking probability, leading the ranking probability to decrease and move further to the left, again until a stable solution with an arrow is reached (or similarly, if the extra click increases the ranking probability, this will lead to a situation where the clicking probability is above the ranking probability, leading the ranking probability to increase and move further right until it also reaches a stable solution with an arrow).




Figure B.1: Limit clicking probability ($\theta_L^{majority}$) as a function of the limit ranking probability (y) for different values of L and for $\alpha = 1$ (left) and $\alpha = 4$ (right). In both panels, $M = 20$, $\mu = 0.9$, $\gamma = 0.33$, $p = 0.55$, and r_1 is uniform.

$y = y_L^{majority}$ to the two equations just above are decreasing in L in the corresponding ranges. Since at the solutions we have $H_L^{minority}(x_L^{minority}; \alpha, \mu, \gamma, p) = 0$ and $H_L^{majority}(y_L^{majority}; \alpha, \mu, \gamma, p) = 0$, we can write the first derivatives with respect to L as:

$$\begin{aligned} \frac{dx_L^{minority}}{dL} &= -\frac{\partial H_L^{minority}(x_L^{minority}; \alpha, \mu, \gamma, p)/\partial L}{\partial H_L^{minority}(x_L^{minority}; \alpha, \gamma, p)/\partial x} = \frac{\partial \theta_L^{minority}/\partial L}{1 - \partial \theta_L^{minority}/\partial x}, \\ \frac{dy_L^{majority}}{dL} &= -\frac{\partial H_L^{majority}(y_L^{majority}; \alpha, \mu, \gamma, p)/\partial L}{\partial H_L^{majority}(y_L^{majority}; \alpha, \gamma, p)/\partial y} = \frac{\partial \theta_L^{majority}/\partial L}{1 - \partial \theta_L^{majority}/\partial y}. \end{aligned}$$

To see that both are negative, we need to check that $\frac{\partial \theta_L^{minority}}{\partial L} \leq 0$, $\frac{\partial \theta_L^{majority}}{\partial L} \leq 0$ and $\frac{\partial \theta_L^{minority}}{\partial x} < 1$, $\frac{\partial \theta_L^{majority}}{\partial y} < 1$. Straightforward computations yield:

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

which are both clearly non-positive always, and negative for interior values $\alpha > 0$ and $0 < \gamma < 1$ and $0 < x < 1$. Also,

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

Clearly these derivatives are all non-negative. However, we now show that they are strictly less than one or the stable solutions. First, it can be checked that, for any limit clicking probability ($0 \leq x$ (or $y \leq 1$)) and for any other parameter values of the model,

$$0 \leq \frac{\gamma \left(\frac{x}{L}\right)^\alpha}{\gamma \left(\frac{x}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-x}{M-L}\right)^\alpha}, \frac{(1-\gamma) \left(\frac{1-x}{M-L}\right)^\alpha}{\gamma \left(\frac{x}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-x}{M-L}\right)^\alpha}, \frac{\gamma(1-\gamma) \left(\frac{x(1-x)}{L(M-L)}\right)^\alpha}{\left((1-\gamma) \left(\frac{x}{L}\right)^\alpha + \gamma \left(\frac{1-x}{M-L}\right)^\alpha\right)^2} \leq 1,$$

which, as already seen in Figure B.1 above, implies that for values of $\alpha \geq 1$ there may be more than one solution to Equations (B.1) and (B.2). Moreover, some of the solutions may have a slope greater or equal to one. However, as explained above, only those solutions are stable which intersect the functions x (or y) from above. Since both x and y clearly have slope equal to 1 everywhere this implies that both $\theta_L^{minority}$ and $\theta_L^{majority}$ must have a slope less than one at the stable solutions and hence must satisfy $\frac{\partial \theta_L^{minority}}{\partial x} < 1$ and $\frac{\partial \theta_L^{majority}}{\partial y} < 1$. This then implies that the respective solutions $x_L^{minority}$ and $y_L^{majority}$ will satisfy $\frac{\partial \hat{r}_{N,L}}{\partial L} \leq 0$ as well as $\frac{\partial \hat{\rho}_{N,L}}{\partial L} \leq 0$ on the relevant ranges and for N sufficiently large, which will also be negative for interior values of the parameters. This shows AOF for both minority and majority outlets with correct signal. To see that AOF also applies to outlets $J \subset M$ with incorrect signal, notice that, if increasing L , decreases $\hat{\rho}_{N,L}$, then, since the total number of outlets M is fixed and $\hat{\rho}_{N,L} + \hat{\rho}_{N,M \setminus L} = 1$, this readily implies that decreasing $M - L$, increases $\hat{\rho}_{N,M \setminus L}$. Therefore, given the AOF for outlets L with the correct signal and taking $J = M \setminus L$ also shows that AOF applies to minority and majority outlets with incorrect signal. This concludes the proof of Proposition 1. \square

Proof of Corollaries 1 and 2. Again, to avoid duplication, we prove directly Corollary 2 stated in Section 6.1, which includes the case of $\alpha = 1$, stated in Corollary 1, as a special case.

Consider Equation (12). Then, when $\alpha = 0$ it is easy to see that:

$$\rho_{N,L} = \begin{cases} 0 & \text{if } L = 0 \\ (1 - \mu)p + (1 - \mu)(1 - p)(1 - \gamma) + \mu p \gamma & \text{if } 1 \leq L = \frac{M-1}{2} \\ \mu p + \mu(1 - p)(1 - \gamma) + (1 - \mu)p \gamma & \text{if } \frac{M+1}{2} \leq L \leq M - 1 \\ 1 & \text{if } L = M, \end{cases}$$

which can be checked is monotonically increasing in L for $\mu \geq \frac{1}{2}$. Now suppose $\alpha > 0$. To see the non-monotonicity, notice again that at $L = 0$, we have $\rho_{n,L} = 0$, and at $L = M - 1$, we have $\rho_{n,L} = 1$, for all n , so that $\rho_{n,L}$ can only increase at $L = 0, M - 1$. For the case $L = \frac{M-1}{2}$, ($M - L = \frac{M+1}{2}$) note that the difference $\theta_{\frac{M+1}{2}}^{majority}(y; \alpha, \mu, \gamma, p) - \theta_{\frac{M-1}{2}}^{minority}(x; \alpha, \mu, \gamma, p)$ can be written as:

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

where, at the relevant solutions, $x = x_L^{majority}$ and $y = y_{M-L}^{majority}$, we have $y \geq x$, such that the expressions in the second and third lines are both non-negative. And since $\mu > \frac{p}{q} > \frac{1}{2}$, the overall difference $(\theta_{M-L}^{majority} - \theta_L^{minority})$ is non-negative. In all other cases, that is, interior values $L \neq 0, \frac{M-1}{2}, M - 1$, $\rho_{N,L}$ is decreasing in L by Proposition 1, for sufficiently large N , and hence also in the limit, for $\rho_{\infty,L} = \mathcal{P}_L$. \square

Proof of Proposition 2. Also here, to avoid duplication, we prove the proposition directly for the case of $\alpha \geq 0$ mentioned in Section 6.1, which includes the case of $\alpha = 1$, stated in Proposition 2,

as a special case.

From the definition of \mathcal{P} in Equation (9), we have:

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

for any given 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 1. In particular, it suffices to consider the following derivatives for $\theta_L^{minority}$ and $\theta_L^{majority}$ defined respectively in Equations (B.3) and (B.4) above:

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

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

$$\begin{aligned} \frac{\partial \theta_L^{minority}}{\partial \mu} &= -p \left(1 - \frac{\gamma \left(\frac{x}{L}\right)^\alpha}{\gamma \left(\frac{x}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-x}{M-L}\right)^\alpha} \right) - (1-p) \frac{(1-\gamma) \left(\frac{x}{L}\right)^\alpha}{\gamma \left(\frac{1-x}{M-L}\right)^\alpha + (1-\gamma) \left(\frac{x}{L}\right)^\alpha} \leq 0, \\ \frac{\partial \theta_L^{majority}}{\partial \mu} &= p \left(1 - \frac{\gamma \left(\frac{y}{L}\right)^\alpha}{\gamma \left(\frac{y}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-y}{M-L}\right)^\alpha} \right) + (1-p) \frac{(1-\gamma) \left(\frac{y}{L}\right)^\alpha}{\gamma \left(\frac{1-y}{M-L}\right)^\alpha + (1-\gamma) \left(\frac{y}{L}\right)^\alpha} \geq 0. \end{aligned}$$

It can be further checked that for $\mu q > p > \frac{1}{2}$, the positive effect of when the websites in L form a majority outweighs the negative effect of when they form a minority. To see this, we can rewrite the above derivatives as:

$$\frac{\partial \theta_L^{minority}}{\partial \mu} = -\theta_{L|\mu=1}^{minority} + \theta_{L|\mu=0}^{minority}, \quad \frac{\partial \theta_L^{majority}}{\partial \mu} = \theta_{L|\mu=1}^{majority} - \theta_{L|\mu=0}^{majority}.$$

Moreover, it can be checked that, at the relevant solutions $x = x_L^{minority}$ and $y = y_{M-L}^{majority}$, taking again as the relevant number of outlets with correct signal, L and $M-L$ for the minority and majority case respectively, we have,

$$\theta_{M-L|\mu=1}^{majority}(y) - \theta_{L|\mu=1}^{minority}(x) \geq 0, \quad \theta_{M-L|\mu=0}^{majority}(y) - \theta_{L|\mu=0}^{minority}(x) \leq 0.$$

Together with the above equations, we can obtain the sign of $\frac{\partial \mathcal{P}}{\partial \mu}$ from:

$$\begin{aligned}
& \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} q^L (1-q)^{M-L} \frac{\partial \theta_L^{\text{minority}}}{\partial \mu} + \sum_{L=\frac{M+1}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} \frac{\partial \theta_L^{\text{majority}}}{\partial \mu} \\
&= \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} q^L (1-q)^{M-L} (-\theta_{L|\mu=1}^{\text{minority}} + \theta_{L|\mu=0}^{\text{minority}}) + \sum_{L=\frac{M+1}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} (\theta_{L|\mu=1}^{\text{majority}} - \theta_{L|\mu=0}^{\text{majority}}) \\
&= \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} q^L (1-q)^{M-L} (-\theta_{L|\mu=1}^{\text{minority}} + \theta_{L|\mu=0}^{\text{minority}}) + \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{M-L} q^{M-L} (1-q)^L (\theta_{M-L|\mu=1}^{\text{majority}} - \theta_{M-L|\mu=0}^{\text{majority}}) \\
&\geq \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} q^L (1-q)^{M-L} ((\theta_{M-L|\mu=1}^{\text{majority}} - \theta_{L|\mu=1}^{\text{minority}}) - (\theta_{M-L|\mu=0}^{\text{majority}} - \theta_{L|\mu=0}^{\text{minority}})) \geq 0,
\end{aligned}$$

for $\frac{1}{2} < p < q$, which in turn implies that $\frac{\partial \mathcal{P}}{\partial \mu} \geq 0$.

To see point 3. that \mathcal{P} is decreasing in γ , provided p not too large, μ not too small, notice that:

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

It can be checked that when $p = \frac{1}{2}$, the sum of the solution to Equation (B.3) at L and the solution to Equation (B.4) at $M-L$ sum to one. Thus, if $x = x_L^{\text{minority}}$ solves Equation (B.3) at L , then $y_{M-L}^{\text{majority}} = 1 - x_L^{\text{minority}}$ solves Equation (B.4) at $M-L$. This then implies that, again at $p \approx \frac{1}{2}$,

$$\left. \frac{\partial \theta_{M-L}^{\text{majority}}}{\partial \gamma} \right|_{1-x} \approx -\frac{(1-p)\mu \left(\frac{x(1-x)}{L(M-L)}\right)^\alpha}{\left(\gamma \left(\frac{x}{L}\right)^\alpha + (1-\gamma) \left(\frac{1-x}{M-L}\right)^\alpha\right)^2} + \frac{p(1-\mu) \left(\frac{x(1-x)}{L(M-L)}\right)^\alpha}{\left((1-\gamma) \left(\frac{x}{L}\right)^\alpha + \gamma \left(\frac{1-x}{M-L}\right)^\alpha\right)^2} = -\left. \frac{\partial \theta_L^{\text{minority}}}{\partial \gamma} \right|_x.$$

Summing these up, yields as the total effect on ex ante efficiency,

$$\begin{aligned}
& \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} q^L (1-q)^{M-L} \frac{\partial \theta_L^{\text{minority}}}{\partial \gamma} - \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{M-L} q^{M-L} (1-q)^L \frac{\partial \theta_L^{\text{minority}}}{\partial \gamma} \\
&= \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} (q^L (1-q)^{M-L} - q^{M-L} (1-q)^L) \frac{\partial \theta_L^{\text{minority}}}{\partial \gamma} \leq 0,
\end{aligned}$$

since $q > \frac{1}{2}$, which in turn implies that $\frac{\partial \mathcal{P}}{\partial \gamma} \leq 0$. But because there are always cases $\frac{M+1}{2} \leq L \leq M-1$, where solutions to Equation (B.4) are interior,³⁵ implying $\frac{\partial \theta_L^{\text{majority}}}{\partial \gamma} < 0$ for some L , summing up actually implies that $\frac{\partial \mathcal{P}}{\partial \gamma} < 0$. This in turn allows us to extend the decreasing effect of γ on \mathcal{P} to a neighborhood $[\bar{\mu}, 1] \times [\frac{1}{2}, \bar{p}]$, for some $\bar{\mu} < 1$, $\bar{p} > \frac{1}{2}$, that exist given the continuity (in fact, linearity) of the derivatives $\frac{\partial \theta_L^{\text{minority}}}{\partial \gamma}$ and $\frac{\partial \theta_L^{\text{majority}}}{\partial \gamma}$ in both μ and p .

The fact that \mathcal{P} is both increasing and decreasing in q is easy to check by means of examples. As

³⁵For example, solutions always satisfy $1 - \mu < x_L^{\text{minority}}, y_L^{\text{majority}} < \mu$, and are interior whenever $\mu \neq 1$.

illustrated in Figure 2, \mathcal{P} is increasing (decreasing) in q for low (high) values of q . \square

Proof of Proposition 3 Assume first that $\mu = 1$. Using the proof of Proposition 1, it can be checked that, for $\alpha = 0$, the solutions to Equations (B.3) and (B.4) take the form, respectively, $x_{L|\alpha=0}^{minority} = \gamma p$, $y_{L|\alpha=0}^{majority} = 1 - \gamma(1 - p)$, whereas, for $\alpha = 1$, they take the form:

$$x_{L|\alpha=1}^{minority} = \begin{cases} \frac{(1-\gamma)L - \gamma p(M-L)}{\gamma^{M-L}} & \text{if } L < \frac{\gamma p M}{1 - (1-\gamma)p} \\ 0 & \text{if } L \geq \frac{\gamma p M}{1 - (1-\gamma)p} \end{cases}, \quad y_{L|\alpha=1}^{majority} = \begin{cases} 1 & \text{if } L \leq \frac{(1-\gamma)M}{1-\gamma p} \\ \frac{\gamma p L}{L - (1-\gamma)M} & \text{if } L > \frac{(1-\gamma)M}{1-\gamma p} \end{cases}.$$

It is easy to see that there exists $\bar{\gamma} (= \frac{1}{p + (1-p)M}) > 0$ such that, for $\gamma \in [0, \bar{\gamma}]$, the solution for $\alpha = 1$ is always $y_{L|\alpha=1}^{majority} = 1$, for $\frac{M+1}{2} \leq L \leq M - 1$, which is greater than the corresponding solution for $\alpha = 0$, ($y_{L|\alpha=0}^{majority} = 1 - \gamma(1 - p)$), for any $\frac{M+1}{2} \leq L \leq M - 1$. For $\gamma \in [0, \bar{\gamma}]$, the corresponding difference between the solutions $x_{L|\alpha=1}^{minority}$ and $x_{L|\alpha=0}^{minority}$ is bounded below by $-\gamma p$. Recall that the solutions coincide with our interim notion of efficiency for the given number of outlets with correct signal, L , so that from the definition of \mathcal{P} in Equation (9), we can write:

$$\begin{aligned} \mathcal{P}' &= 0 + \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} q^L (1-q)^{M-L} x_{L|\alpha=0}^{minority} + \sum_{L=\frac{M+1}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} y_{L|\alpha=0}^{majority} + q^M, \\ \mathcal{P} &= 0 + \sum_{L=1}^{\frac{M-1}{2}} \binom{M}{L} q^L (1-q)^{M-L} x_{L|\alpha=1}^{minority} + \sum_{L=\frac{M+1}{2}}^{M-1} \binom{M}{L} q^L (1-q)^{M-L} y_{L|\alpha=1}^{majority} + q^M, \end{aligned}$$

Since $q \gg p > \frac{1}{2}$, this is enough to imply that, for arbitrarily large q , while $\gamma \in [0, \bar{\gamma}]$, we have, $\mathcal{P} \geq \mathcal{P}'$ and hence $\text{PoR} \geq 0$.

Suppose now $\gamma \geq \bar{\gamma}$. Again, for the given γ , there are two parts to the solution $y_{L|\alpha=1}^{majority}$ for $\alpha = 1$, namely, a part which is 1 and therefore above the solution for $\alpha = 0$ (for $\frac{M+1}{2} \leq L \leq \frac{(1-\gamma)M}{1-\gamma p}$) and a part which is below the solution for $\alpha = 0$ (for $\frac{(1-\gamma)M}{1-\gamma p} < L \leq M - 1$). In particular, there exists a level \bar{q} such that, when $q \leq \bar{q}$, then the cases where $L \leq \frac{(1-\gamma)M}{1-\gamma p}$ and hence where the solution $y_{L|\alpha=1}^{majority}$ for $\alpha = 1$ is above the solution $y_{L|\alpha=0}^{majority}$ for $\alpha = 0$ will obtain sufficiently large weight, such that ex ante efficiency for $\alpha = 1$ (\mathcal{P}) is greater or equal to ex ante efficiency for $\alpha = 0$ (\mathcal{P}'), that is, such that $\text{PoR} \geq 0$. At the same time, when $q > \bar{q}$, then the cases where $L > \frac{(1-\gamma)M}{1-\gamma p}$ and hence where the solution $y_{L|\alpha=1}^{majority}$ for $\alpha = 1$ is below the solution $y_{L|\alpha=0}^{majority}$ for $\alpha = 0$ will obtain sufficiently large weight, such that ex ante efficiency for $\alpha = 1$ (\mathcal{P}) is less or equal to ex ante efficiency for $\alpha = 0$ (\mathcal{P}'), that is, such that $\text{PoR} \leq 0$. This gives the function ϕ , which is equal to 1 on $[0, \bar{\gamma}]$ and which is decreasing in γ on $[\bar{\gamma}, 1]$, since the set of interim realizations, where the solution $y_{L|\alpha=1}^{majority}$ for $\alpha = 1$ is above the solution $y_{L|\alpha=0}^{majority}$, is determined by the cutoff $\frac{(1-\gamma)M}{1-\gamma p}$, which is decreasing in γ . See Figure B.2 (left panel) for an illustration of the case with $\mu = 1$.

Finally, the fact that there exists $\bar{\mu} < 1$, such that the above holds for $\mu \in [\bar{\mu}, 1]$, follows essentially from the fact that the functions $\theta_L^{minority}$ and $\theta_L^{majority}$ defining the Equations (B.3) and (B.4) are linear in μ . This implies that the solutions are arbitrarily close to the ones computed for $\mu = 1$. More specifically, it can be checked that, for general μ , the solutions to Equations (B.3) and (B.4) for $\alpha = 0$ take the form, $x_{L|\alpha=0}^{minority} = (1 - \mu)(1 - \gamma) + \gamma p$, $y_{L|\alpha=0}^{majority} = \mu(1 - \gamma) + \gamma p$. At the same time, while the solutions for $\alpha = 1$ no longer have a simple closed form, it can be checked that, for




Figure B.2: Interim efficiency ($\mathcal{P}_L, \mathcal{P}'_L$) for $\alpha = 0$ (black) and $\alpha = 1$ (red dashed) as a function of L for $\mu = 1$ (left) and $\mu = 0.9$ (right). In both panels, $M = 20$, $\gamma = 0.33$, $p = 0.55$, and r_1 is uniform. The left panel also shows the density function for the binomial distribution for $q = 0.7 < \phi(\gamma)$ (light gray dashed) at which $\text{PoR} > 0$, while the right panel shows the density function for the binomial distribution for $q = 0.9 > \phi(\gamma)$ at which $\text{PoR} < 0$.

given L , they are above the corresponding solutions for $\alpha = 0$, for $L < ((1 - \mu)(1 - \gamma) + \gamma p)M$, if L is minority, and for $L < (\mu(1 - \gamma) + \gamma p)M$, if L is majority. Again, checking for the levels of ex ante efficiency, this allows to compute decreasing cutoff levels for q as a function of γ, \bar{q} , such that for $q \geq \bar{q}$, $\text{PoR} \leq 0$, (provided γ is not too small), while for $q \geq \bar{q}$, $\text{PoR} \geq 0$. See Figure B.2 (right panel) for an illustration of the case with $\mu < 1$. \square

Proof of Proposition 4. For simplicity, we consider the case, where one half of the population has γ_A and the other half has γ_B , and where $\gamma_A \neq \gamma_B$. Recall the equation defining the ranking probabilities in the case of personalization (Equations (10) and (11)), for any t and m , and for $\ell = A, B$:

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

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:

$$\begin{aligned} \hat{r}_{t,m}^A &= \hat{r}_{t-1,m}^A + \frac{1}{1 + \kappa_t} (\hat{\rho}_{t-1,m}^A - \hat{r}_{t-1,m}^A) + \frac{1 - \lambda}{1 - \lambda + \kappa_t} (\hat{\rho}_{t-1,m}^B - \hat{r}_{t-1,m}^A) \\ \hat{r}_{t,m}^B &= \hat{r}_{t-1,m}^B + \frac{1}{1 + \kappa_t} (\hat{\rho}_{t-1,m}^B - \hat{r}_{t-1,m}^B) + \frac{1 - \lambda}{1 - \lambda + \kappa_t} (\hat{\rho}_{t-1,m}^A - \hat{r}_{t-1,m}^B), \end{aligned} \quad (\text{B.5})$$

where, following Equation (A.1), we can write, for $\ell = A, B$:

$$\begin{aligned} \hat{\rho}_{t-1,m}^\ell &= \mathbb{E}[\rho_{t-1,m}^\ell] = p\mu \frac{(\hat{r}_{t-1,m}^\ell)^\alpha \cdot \hat{v}_m^{*00}(\gamma_\ell)}{\sum_{m'} (\hat{r}_{t-1,m'}^\ell)^\alpha \cdot \hat{v}_{m'}^{*00}(\gamma_\ell)} + p(1 - \mu) \frac{(\hat{r}_{t-1,m}^\ell)^\alpha \cdot \hat{v}_m^{*01}(\gamma_\ell)}{\sum_{m'} (\hat{r}_{t-1,m'}^\ell)^\alpha \cdot \hat{v}_{m'}^{*01}(\gamma_\ell)} \\ &\quad + (1 - p)\mu \frac{(\hat{r}_{t-1,m}^\ell)^\alpha \cdot \hat{v}_m^{*10}(\gamma_\ell)}{\sum_{m'} (\hat{r}_{t-1,m'}^\ell)^\alpha \cdot \hat{v}_{m'}^{*10}(\gamma_\ell)} + (1 - p)(1 - \mu) \frac{(\hat{r}_{t-1,m}^\ell)^\alpha \cdot \hat{v}_m^{*11}(\gamma_\ell)}{\sum_{m'} (\hat{r}_{t-1,m'}^\ell)^\alpha \cdot \hat{v}_{m'}^{*11}(\gamma_\ell)}. \end{aligned}$$

Taking the limit $\kappa_t \rightarrow \infty$ in the equation system (B.5), yields the equations determining the limit

clicking probabilities with personalization, which take the form:

$$\frac{1}{2-\lambda}\widehat{\rho}_{t,m}^A + \frac{1-\lambda}{2-\lambda}\widehat{\rho}_{t,m}^B - r_{t,m}^A = 0 \quad \text{and} \quad \frac{1}{2-\lambda}\widehat{\rho}_{t,m}^B + \frac{1-\lambda}{2-\lambda}\widehat{\rho}_{t,m}^A - r_{t,m}^B = 0.$$

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:

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

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\mu \cdot (x_m)^\alpha \cdot \widehat{v}_m^{*00}(\gamma_A)}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{v}_{m'}^{*00}(\gamma_A)} + \frac{p(1-\mu) \cdot (x_m)^\alpha \cdot \widehat{v}_m^{*01}(\gamma_A)}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{v}_{m'}^{*01}(\gamma_A)} + \frac{(1-p)\mu \cdot (x_m)^\alpha \cdot \widehat{v}_m^{*10}(\gamma_A)}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{v}_{m'}^{*10}(\gamma_A)} + \frac{(1-p)(1-\mu) \cdot (x_m)^\alpha \cdot \widehat{v}_m^{*11}(\gamma_A)}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{v}_{m'}^{*11}(\gamma_A)} \right) \\ &\quad + \frac{1-\lambda}{2-\lambda} \left(\frac{p\mu \cdot (y_m)^\alpha \cdot \widehat{v}_m^{*00}(\gamma_B)}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{v}_{m'}^{*00}(\gamma_B)} + \frac{p(1-\mu) \cdot (y_m)^\alpha \cdot \widehat{v}_m^{*01}(\gamma_B)}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{v}_{m'}^{*01}(\gamma_B)} + \frac{(1-p)\mu \cdot (y_m)^\alpha \cdot \widehat{v}_m^{*10}(\gamma_B)}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{v}_{m'}^{*10}(\gamma_B)} + \frac{(1-p)(1-\mu) \cdot (y_m)^\alpha \cdot \widehat{v}_m^{*11}(\gamma_B)}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{v}_{m'}^{*11}(\gamma_B)} \right), \\ \theta_m^B(x, y) &= \frac{1-\lambda}{2-\lambda} \left(\frac{p\mu \cdot (x_m)^\alpha \cdot \widehat{v}_m^{*00}(\gamma_A)}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{v}_{m'}^{*00}(\gamma_A)} + \frac{p(1-\mu) \cdot (x_m)^\alpha \cdot \widehat{v}_m^{*01}(\gamma_A)}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{v}_{m'}^{*01}(\gamma_A)} + \frac{(1-p)\mu \cdot (x_m)^\alpha \cdot \widehat{v}_m^{*10}(\gamma_A)}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{v}_{m'}^{*10}(\gamma_A)} + \frac{(1-p)(1-\mu) \cdot (x_m)^\alpha \cdot \widehat{v}_m^{*11}(\gamma_A)}{\sum_{m'} (x_{m'})^\alpha \cdot \widehat{v}_{m'}^{*11}(\gamma_A)} \right) \\ &\quad + \frac{1}{2-\lambda} \left(\frac{p\mu \cdot (y_m)^\alpha \cdot \widehat{v}_m^{*00}(\gamma_B)}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{v}_{m'}^{*00}(\gamma_B)} + \frac{p(1-\mu) \cdot (y_m)^\alpha \cdot \widehat{v}_m^{*01}(\gamma_B)}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{v}_{m'}^{*01}(\gamma_B)} + \frac{(1-p)\mu \cdot (y_m)^\alpha \cdot \widehat{v}_m^{*10}(\gamma_B)}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{v}_{m'}^{*10}(\gamma_B)} + \frac{(1-p)(1-\mu) \cdot (y_m)^\alpha \cdot \widehat{v}_m^{*11}(\gamma_B)}{\sum_{m'} (y_{m'})^\alpha \cdot \widehat{v}_{m'}^{*11}(\gamma_B)} \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 if $\lambda = 0$, then it is as if there were a single group with the same (common) ranking algorithm and where the individuals' preference for like-minded news is $\gamma = \frac{\gamma_A + \gamma_B}{2}$. As λ increases, then the rankings of the two groups drift apart and it is as if group A had parameter $\frac{1}{2-\lambda}\gamma_A + \frac{1-\lambda}{2-\lambda}\gamma_B$ and group B had parameter $\frac{1-\lambda}{2-\lambda}\gamma_A + \frac{1}{2-\lambda}\gamma_B$ until, when $\lambda = 1$, it is as if there were two separate rankings of two groups with accuracy γ_A and γ_B respectively. Since $\gamma_A \neq \gamma_B$ the difference in the levels of preference for like-minded news of the two groups $(\frac{\lambda}{2-\lambda}(\gamma_A - \gamma_B))$ 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 5. 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 (B.3) and (B.4) in the proof of Proposition 1, where we recall again that $x_L^{minority} = \widehat{\rho}_{\infty, L} = \mathcal{P}_L^\lambda$ when websites in L have minority signal and $y_L^{majority} = \widehat{\rho}_{\infty, L} = \mathcal{P}_L^\lambda$ 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 $1 < L < M$. As in the proof of Proposition 3, we have that, for $\alpha = 1$ and $\mu = 1$, the solutions to Equations (B.3) and (B.4) take the form, respectively:

$$x_L^{minority} = \begin{cases} \frac{(1-\gamma)L - \gamma p(M-L)}{\gamma^{M-L}} & \text{if } L < \frac{\gamma p M}{1 - (1-\gamma)p} \\ 0 & \text{if } L \geq \frac{\gamma p M}{1 - (1-\gamma)p} \end{cases}, \quad y_L^{majority} = \begin{cases} 1 & \text{if } L \leq \frac{(1-\gamma)M}{1-\gamma p} \\ \frac{\gamma p L}{L - (1-\gamma)M} & \text{if } L > \frac{(1-\gamma)M}{1-\gamma p} \end{cases}.$$

Fix a realization of \mathcal{E} , say, parametrized by L . From the proof of Proposition 4, we have that when $\lambda = 0$, it is as if there were a single group with a common ranking algorithm, with $\gamma = \frac{\gamma_A + \gamma_B}{2}$. As λ increases, then the rankings of the two groups drift apart and are as if group A had preference




Figure B.3: Interim efficiency with no personalization $\mathcal{P}_L^{\lambda=0}$ (black) and with full personalization $\mathcal{P}_L^{\lambda=1}$ (green dashed) as a function of L for $\mu = 1$ (left) and $\mu = 0.9$ (right). In both panels, $M = 20$, $\gamma_A = 0$, $\gamma_B = 0.66$, $p = 0.55$, and r_1 is uniform. Both panels also show (gray dashed) the function $\frac{L}{M}$ as comparison benchmark as well as the density function for the binomial distribution for $q = 0.7 < \phi(\gamma)$ at which $\text{PeR} > 0$.

for like-minded news $\frac{1-\lambda}{2-\lambda}\gamma_A + \frac{1-\lambda}{2-\lambda}\gamma_B$ and group B had $\frac{1-\lambda}{2-\lambda}\gamma_A + \frac{1}{2-\lambda}\gamma_B$, until, when $\lambda = 1$, it is as if there were two separate rankings of two groups with accuracy γ_A and γ_B , respectively. As a result interim efficiency can be written as the average of the interim efficiency of two groups, one with $\gamma_A^\lambda = \frac{1-\lambda}{2-\lambda}\gamma_A + \frac{1-\lambda}{2-\lambda}\gamma_B$ and another with $\gamma_B^\lambda = \frac{1-\lambda}{2-\lambda}\gamma_A + \frac{1}{2-\lambda}\gamma_B$. 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 $\frac{\gamma_A + \gamma_B}{2}$ when $\lambda = 0$ to being γ_A and γ_B respectively when $\lambda = 1$). Since we can always take γ_A, γ_B to be, respectively, $\gamma_A^\lambda, \gamma_B^\lambda$, it suffices to consider directly the case of $\lambda = 1$. In this case, the interim efficiency levels can be computed from the solutions evaluated at the corresponding levels of γ :

$$\mathcal{P}_L^{\lambda=0} = \left(x_L^{\text{minority}} \left(\frac{\gamma_A + \gamma_B}{2} \right), y_L^{\text{majority}} \left(\frac{\gamma_A + \gamma_B}{2} \right) \right),$$

for the non-personalized case ($\lambda = 0$), and from:

$$\mathcal{P}_L^{\lambda=1} = \left(\frac{x_L^{\text{minority}}(\gamma_A) + x_L^{\text{minority}}(\gamma_B)}{2}, \frac{y_L^{\text{majority}}(\gamma_A) + y_L^{\text{majority}}(\gamma_B)}{2} \right),$$

for the fully personalized case ($\lambda = 1$). Since the reasoning is very similar to that of the proof of Proposition 3, we focus on the case where the outlets with correct signal are a majority. For $\lambda = 0$ we have, as the level of interim efficiency, directly $y_L^{\text{majority}} \left(\frac{\gamma_A + \gamma_B}{2} \right)$ from above, while for $\lambda = 1$, recalling $0 \leq \gamma_A < \gamma_B \leq 1$, and since the cutoff for L ($\frac{\gamma p M}{1 - (1-\gamma)p}$) is decreasing in γ , we can write this as:

$$\frac{y_L^{\text{majority}}(\gamma_A) + y_L^{\text{majority}}(\gamma_B)}{2} = \begin{cases} 1 & \text{if } \frac{M+1}{2} \leq L \leq \frac{(1-\gamma_B)M}{1-\gamma_{BP}} \\ \frac{1}{2} + \frac{\gamma_{BP}L}{2(L-(1-\gamma_B)M)} & \text{if } \frac{(1-\gamma_B)M}{1-\gamma_{BP}} < L \leq \frac{(1-\gamma_A)M}{1-\gamma_{AP}} \\ \frac{\gamma_{AP}L}{2(L-(1-\gamma_A)M)} + \frac{\gamma_{BP}L}{2(L-(1-\gamma_B)M)} & \text{if } \frac{(1-\gamma_A)M}{1-\gamma_{AP}} < L \leq M-1 \end{cases}.$$

Tedious calculations show that, when $y_L^{\text{majority}} \left(\frac{\gamma_A + \gamma_B}{2} \right) \geq \frac{L}{M}$, we also have $y_L^{\text{majority}} \left(\frac{\gamma_A + \gamma_B}{2} \right) \geq \frac{y_L^{\text{majority}}(\gamma_A) + y_L^{\text{majority}}(\gamma_B)}{2}$, and hence $\mathcal{P}_L^{\lambda=0} \geq \mathcal{P}_L^{\lambda=1}$.³⁶ Using the same reasoning as in the proof

³⁶To get more intuition for the proof, notice that $\mathcal{P}_L^{\lambda=0}$ for majority values of L is a concave function of γ , when solutions satisfy $y \geq \frac{L}{M}$, and becomes convex when $y \leq \frac{L}{M}$. Interim efficiency being concave implies ex ante efficiency

of Proposition 3, we can show that y (or also x) $\geq \frac{L}{M}$ occurs, when μ is sufficiently large and γ sufficiently small, and when q is below a given threshold $\bar{q}(\gamma)$. It can further be shown that the reverse inequality holds for the same values of μ and γ if q is above the threshold $\bar{q}(\gamma)$. Finally, a similar argument as in the proof of Proposition 3 allows to extend to the case where $\mu > \bar{\mu}$ for some $\bar{\mu} < 1$. Figure B.3 illustrates the cases $\mu = 1$ (left panel) and $\mu < 1$ (right panel). \square

Proof of Corollary 2. This is proved above together with Corollary 1. \square

Proof of Proposition 6 and Proposition 7. We prove directly the case with $\gamma \geq 0$. The expected ranking probabilities are then given by:

$$\begin{aligned} \hat{r}_{n,m} &= \nu \hat{r}_{n-1,m} + (1-\nu) \hat{\rho}_{n-1,m} \\ &= \nu \hat{r}_{n-1,m} + (1-\nu) \left(\frac{p\mu(\hat{r}_{t-1,m})^\alpha \cdot \hat{v}_m^{*00}}{\sum_{m'}(\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*00}} + \frac{p(1-\mu)(\hat{r}_{t-1,m})^\alpha \cdot \hat{v}_m^{*01}}{\sum_{m'}(\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*01}} + \right. \\ &\quad \left. + \frac{(1-p)\mu(\hat{r}_{t-1,m})^\alpha \cdot \hat{v}_m^{*10}}{\sum_{m'}(\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*10}} + \frac{(1-p)(1-\mu)(\hat{r}_{t-1,m})^\alpha \cdot \hat{v}_m^{*11}}{\sum_{m'}(\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*11}} \right) \\ &= \nu \hat{r}_{n-1,m} + (1-\nu) \left(\frac{p\mu \cdot \hat{v}_m^{*00}}{\sum_{m'}(\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*00}} + \dots + \frac{(1-p)(1-\mu) \cdot \hat{v}_m^{*11}}{\sum_{m'}(\hat{r}_{t-1,m'})^\alpha \cdot \hat{v}_{m'}^{*11}} \right) \cdot (\hat{r}_{n-1,m})^\alpha, \end{aligned}$$

where \hat{v}_m^{*00} , \hat{v}_m^{*01} , \hat{v}_m^{*10} and \hat{v}_m^{*11} 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 $x \equiv \hat{r}_{n-1,m}$ and $y \equiv \hat{r}_{n-1,m'}$. 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 $(\hat{r}_{n-1,m})^\alpha$ and $\hat{r}_{n-1,m}$, say, a and b respectively, where $a, b > 0$. Hence we can write:

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

But then it follows that:

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

and similarly

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

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

is concave in γ such that ex ante efficiency of an average $\gamma = \frac{\gamma_A + \gamma_B}{2}$ will be above the average ex ante efficiency of γ_A and γ_B (which in turn implies $PeR \leq 0$, provided $q < \bar{q}$). On the other hand, the opposite holds ($PeR \geq 0$), when the function becomes convex. This shows why larger levels of q ($> \bar{q}$) can reverse the result.

$\frac{a}{1-\nu} > 0$, 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 $\alpha = 1$. □

ONLINE APPENDIX

The Online Appendix presents examples that provide numerical and graphical illustrations of some results or of extensions of results presented in the paper.

A. Advantage of the Fewer

First, we provide a numerical illustration of the *amplification effect* generated by the *AOF*. Then, we show that the *AOF* is a rather general phenomenon that holds even outside the assumptions of the model. In particular, we consider three cases not covered by the basic model studied in the paper, namely, (1) a case with uninformative signals, (2) a case where the ranking is an ordinal ranked list and not a probability distribution over the outlets, (3) a case where the total number of outlets is not fixed such that, for example, two outlets merging and thus reducing the total number of outlets. Finally, we illustrate the stochastic dynamic behind the *AOF*.

Example A1. *AOF Amplification effect.* As discussed in Section 3, the *AOF* generates an *amplification effect* whereby a reduction in the number of websites reporting a given signal increases their overall share of traffic. Such an effect can be sizeable. Consider, for example, a search environment where $M = 20$, r_1 uniform and where $p = 0.55$ and $\mu = 0.9$. Figure A.1 illustrates the *amplification effect* due to a reduction in the number of outlets reporting a correct signal, as a function of the preference for like-minded news. As the graph shows, for the case of minority outlets (e.g., a reduction in L from 2 to 1), the *amplification effect* may almost double the probability of clicking on a minority outlet reporting a correct signal. In the case of majority outlets (e.g., a reduction in L from 16 to 15) the effect is smaller, but it may still boost the overall traffic of such outlets by up to 15%.

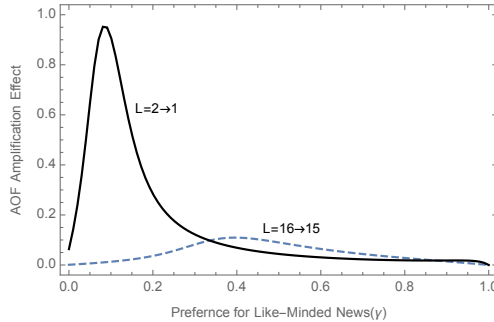


Figure A.1: Amplification effect of *AOF* (as a function of the preference for like-minded news, γ) in terms of limit clicking probability of minority websites reporting a correct signal (black) when L is reduced from 2 to 1 (black) and of majority websites reporting a correct signal (dashed) when L is reduced from 16 to 15. In both panels, $M = 20$, $p = 0.55$, $\mu = 0.9$, and r_1 is uniform.

Example A2. AOF with uninformative signals. Consider a search environment where $M = 20$, r_1 uniform and where $p = \mu = \frac{1}{2}$. Then, as Figure A.2 shows, the limit clicking probability of choosing an outlet reporting a correct signal is monotonically decreasing in the number of websites with a correct signal (for $1 \leq L < M - 1$). This also shows that the AOF effect can be generated with just flows of types of individuals, ones that are committed to one type of content and therefore actively search for those outlets, wherever they appear in the ranking, ones that are committed to the other type of content, and ones that are willing to trade off one content type with the other. As discussed in the text, it is the latter that generate the amplification effect.




Figure A.2: *Interim* efficiency (\mathcal{P}_L) as a function of the number of websites with correct signal (L) for $\gamma = 0.33$ (black line) and other values of γ (dashed lines) with uninformative signals $p = 0.5$, $\mu = 0.5$. Also, $M = 20$, and r_1 is uniform.

Example A3. AOF with ordinal ranking and pure clicking realizations. Consider a search environment, where the ranking at each time is a list from 1 to 20, and where the rank of an outlet $r_{t,m} \in \{1, \dots, M\}$ depends on the number of clicks received up to t , such that the outlet with the greatest number of clicks has $r_{t,m} = 1$, the one with the second greatest number of clicks has $r_{t,m} = 2$ and so on. Suppose also that each individual chooses a single outlet according to the vector of choice probabilities ρ_n , thereby contributing a single click rather than a probability vector (as implicitly assumed in Equation (5)). In this case, we need to change the weighting function used to describe the individual stochastic choice. We use the function $\beta^{(M-r_{t,m})}$, where $\beta \geq 1$ again calibrates attention bias in a way that being one position higher corresponds to receiving β times as much probability of being clicked. Otherwise, we use the same multinomial probabilities of choosing a website given by:

$$\rho_{t,m} = \frac{\beta^{(M-r_{t,m})} v_{t,m}^*}{\sum_{m' \in M} \beta^{(M-r_{t,m'})} v_{t,m'}^*}.$$

Suppose also that r_1 is a list that has all minority websites at the bottom, and where $M = 20$, $\beta = 1.5$, $p = 0.55$, $\mu = 0.9$. Then, the following figure shows that the limit clicking probability of

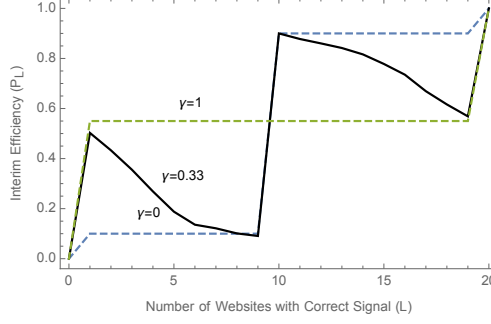


Figure A.3: *Interim* efficiency (\mathcal{P}_L) as a function of the number of websites with correct signal (L) for $\gamma = 0.33$ (black line) and other values of γ (dashed lines) with ordinal ranking and pure clicking realizations. $M = 20$, $p = 0.55$, $\mu = 0.9$, $\beta = 1.5$ and r_1 has all outlets L with correct signal ranked at the bottom L positions.

choosing an outlet reporting a correct signal is again decreasing in the number of websites with a correct signal for $L \neq 1, \frac{M-1}{2}, M-1$. Notice that it is very similar to Figure 1 obtained with the probabilistic ranking and clicking.

Example A4. AOF with merging of outlets. Consider a search environment and consider interim realizations where the number of outlets with the correct signal varies, but where, unlike the case with a fixed total number of outlets M characterizing the basic framework of the paper, the total number of outlets can vary. In particular, we want to allow two outlets with the same signal to merge such that, after they merge, the total number of outlets decreases by one. To give an idea of what happens in this case, we show the *AOF* effect in an environment, where there is a fixed number of outlets with the wrong signal, say $J = 10$, and where the number of outlets with the correct signal varies from $L = 0$ to $L = 20$. In this case the total number of outlets varies from $M = 10$ to $M = 30$. As is illustrated in Figure A.4. The effect is robust to considering this way of changing the number of outlets with a given signal. (Note that we could also consider the case where L is fixed and the number of outlets with the wrong signal varies.)




Figure A.4: *Interim* efficiency (\mathcal{P}_L) as a function of the number of websites with correct signal (L) for $\gamma = 0.33$ (black line) and other values of γ (dashed lines) when the number of outlets with incorrect signal is fixed at $J = 10$. Also, $p = 0.55$, $\mu = 0.9$ and r_1 has all outlets L with correct signal ranked at the bottom L positions.

Example A5. AOF and stochastic dynamics. Consider the search environment with $M = 20$, $p = 0.55$, $\mu = 1$, and $\gamma > 0$, and assume that the initial ranking r_1 is uniform. Suppose there are three websites reporting a website-minority signal. Two cases may arise. The signal is the correct one ($L = 3$) or is the incorrect one ($M - L = M - K = 3$). In either case, the three minority websites gain from the fact that there are few websites that report the same signal. As a result they can attract a significant share of traffic even in the case where they report the incorrect signal. The share depends on γ and can reach over 50% total clicking probability for γ sufficiently large. More specifically, for small values of γ it is close to zero for $L = 3$ or $M - L = 3$. As γ increases, more ranking probability is put on the minority websites, until it reaches a maximum total ranking probability for values of γ sufficiently close to 1.




Figure A.5: Left panel: Total ranking probability on three minority websites with correct signal ($L = M - K = 3$) (black line); the dashed lines represent the mean dynamics trajectories for $L = 1$ (orange), $L = 3$ (green) and $L = 5$ (red). Right panel: Total ranking probability on three minority websites with wrong signal ($M - L = M - K = 3$) (black line); the dashed lines represent the mean dynamics trajectories for $M - L = 1$ (orange), $M - L = 3$ (green) and $M - L = 5$ (red) (the mean dynamics is discussed in Appendix A). In both cases $\gamma = 0.5$ and the initial ranking is uniform.

Decreasing the number of minority websites to one or two ($L = 1, 2$ or $M - L = 1, 2$) *increases* the total ranking probability of these websites as compared to the case of three. Similarly, increasing the number of minority websites to four or five or more, *decreases* the total ranking probability of these websites as compared to the case of three. That is, there is an *advantage of the fewer*. Figure A.5 shows the evolution of the total clicking probability for the cases of $L = 3$ (left panel) and $M - L = 3$ (right panel). At the same time it shows the expected total ranking probabilities for $L = 3$ ($M - L = 3$) and also for $L = 1$ ($M - L = 1$) and for $L = 5$ ($M - L = 5$). The figure illustrates that decreasing the number of minority websites *increases* the total ranking probability of those websites regardless of whether they are reporting the correct signal or not.

B. Rich get richer

Consider now a search environment where individuals do not have any preference for like-minded news ($\gamma = 0$) and where there is no personalization ($\lambda = 0$). The following numerical example shows how the attention bias (α) and the initial ranking (r_1) may create a rich-get-richer dynamic.

Example B1. Initial ranking, attention bias and rich get richer. Consider an environment with $(N, \alpha, \gamma) = (1000, \alpha, \gamma)$; $M = 20$; $(\kappa, \lambda) = (100, 0)$. For now, assume $\gamma = 0$, $\mu = 1$, and consider an interim realization with $L = 15$ and a uniform initial ranking r_1 . Then, agent n 's ranking-free website choice is:

$$v_{n,m}^* = \begin{cases} \frac{1}{15} & \text{if } y_m = y_K \\ 0 & \text{else,} \end{cases}$$

which immediately implies $\rho_{n,m} = v_{n,m}^*$, for $n \in N$, $m \in M$. This means that all individuals always only read websites with the website-majority signal $y_L = y_K$, and access any one of them with the same probability $\frac{1}{15}$ so that, with $\gamma = 0$ the website-majority signal prevails. Moreover, since the initial ranking is uniform, the traffic is evenly distributed across websites carrying such a signal.

Non-uniform initial ranking. Suppose now that the initial ranking is given by:

$$r_1 \approx (0.06, 0.059, \dots, 0.041, 0.04),$$

that is, the websites are ranked with increments of $\frac{1}{950} \approx 0.001$ starting from 0.04 going up to 0.06. Again, since $\gamma = 0$, individuals always choose one of the fifteen websites reporting the website-majority signal $y_L = y_K$. At the same time, the presence of a non-uniform initial ranking provides an advantage to the websites with a higher initial ranking. In particular, as illustrated in Figure B.1, increasing the attention bias parameter α , makes for even larger effects of the websites' initial ranking.

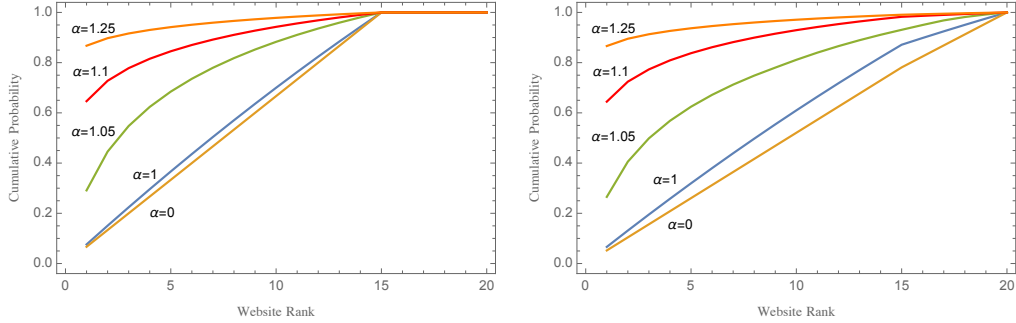


Figure B.1: Cumulative probability of clicks per website for attention bias $\alpha = 0, 1, 1.05, 1.1, 1.25$. The graph is drawn for an initial ranking $r_1 = (0.06, 0.059, \dots, 0.041, 0.04)$ and for $\gamma = 0$ (left panel) and $\gamma = 0.5$ (right panel).

Comparing this pattern with the one arising with a uniform initial ranking, we see that, with a non-uniform initial ranking and $\alpha > 1$, we obtain a *rich-get-richer* dynamic, whereby the ratio of the expected clicking probabilities of two websites $m, m' \in K$, with the same website-majority signal and 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. When $0 \leq \alpha \leq 1$ the clicking probabilities tend towards uniform ranking over the websites in K (when $\alpha < 1$) and a normalization of the initial ranking again over the websites in K (when $\alpha = 1$). A similar pattern applies when $\gamma > 0$ (right panel). \square

This example shows that the evolution of a website’s ranking based on its “popularity,” interacted with a sufficiently large attention bias ($\alpha > 1$) of the individuals exhibits a *rich-get-richer* dynamic, which in turn implies a tendency towards concentration of website traffic at the top. 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 attention bias.