Agenda Chasing and Contests Among News Providers

Zsolt Katona, Jonathan Knee and Miklos Sarvary[1]

May, 2013

[1]Zsolt Katona is Assistant Professor of Marketing at the Haas School of Business at UC Berkeley, Jonathan Knee is Senior Managing Director at Evercore Partners and Adjunct Professor of Finance and Economics at Columbia Business School and Miklos Sarvary is Professor of Marketing at Columbia Business School. The authors would like to thank Francesco Marconi for helpful examples and insights. Preliminary draft. Comments welcome.

Abstract

We model competition among news providers as a contest where each firm chooses to publish on a topic from a large pool of topics with different prior success probabilities. If a topic is successful, firms that chose to publish on it share a fixed reward. We explore how increased competition (as measured by the number of firms and/or the share-structure of the reward) and the prior distribution of topics affect the diversity of published news. We relate our findings to current trends in news media, characterized by lower barriers to entry and the increased use of sophisticated technologies to identify successful topics from the large amount of, both professional and user-generated content available on the Internet. We show that the contest nature of competition tends to lead to a broader set of published media themes with a higher representation for marginal topics. The breadth of topics increases the more topics follow a "fat-tail" prior distribution and the more *a priori* popular topics' success are correlated. It also increases with the number of competing firms but only if the share of the reward in the contest dissipates rapidly. We also explore the effect of asymmetry on competition. First, we assume that some firms have a 'brand', i.e. a capability to attract a loyal audience. We show that branded publishers are more likely to choose topics with high prior success probabilities, while unbranded publishers tend to choose *a priori* 'unlikely' topics. Second, we assume that some firms have better forecasting capability for the topics' success. Surprisingly, in this case, less informed firms choose topics in a conservative way (i.e. publish topics with the highest prior probabilities). When many firms reporting on the same topic increases the topic's rate of success, marginal topics may emerge but only if the contest is not too competitive and competing firms are neither too few nor too numerous. These findings are related to current trends in the news media industry.

Keywords: agenda setting, game theory, media competition.

1 Introduction

The Internet has dramatically changed, not only the way people acquire information and news but also the general topics and themes that emerge as central subjects of interest to the public. While traditional mainstream media (the press or broadcast media) still have a strong role in so-called "agenda setting", increasingly, user-generated content represents serious competition to these outlets as seemingly marginal topics - not to say, trivia - manage to grab disproportionate attention from the public, often with the help of new (online) publishers. This in turn has consequences for reader/viewership and ultimately for advertising revenues. To illustrate the phenomenon, consider the recent example of the YouTube video Gangnam Style. A single by Jae-Sang Park of PSY that was quickly becoming famous in Korea, it was first identified and popularized by Gawker, an online publishing outlet specialized on gossip and quirky urban content.[1] After the Gawker post, in a few month the video became the most viewed clip on YouTube with close to a billion views by the fall of 2012, producing many spinoffs in various countries (e.g. a popular "Mitt Romney Style" version in the U.S.). Mainstream media, including the conservative Fox News, picked up the phenomenon and discussed its significance for contemporary society.[2] While Gangnam Style's success could be explained on the grounds of fundamental artistic value it is harder to make the case for a similar YouTube phenomenon, "Harlem Shake", which started in February 2013.[3] Thousands of re-interpretations of the original song by Baauer have been made by anonymous users as well as by notable people including journalist Anderson Cooper, The Daily Show Staff and Stephen Colbert. What started as a viral sensation on YouTube soon spread to traditional news outlets being covered

[1]http://gawker.com/5930283/did-this-underground-hip-hop-artist-from-south-korea-just-release-the-best-music-video-of-the-year.

[2]http://www.huffingtonpost.com/2012/11/28/bill-oreilly-gangnam-style_n_2203882.html.

[3]Harlem Shake consists of a genre of videos, each lasting for 31 seconds and featuring an excerpt from the song "Harlem Shake" by electronic musician Baauer. Usually, a video begins with one person (often helmeted or masked) dancing to a beat alone for 15 seconds, surrounded by other people not paying attention or unaware of the dancing individual. When the bass drops, the video cuts to the entire crowd doing a strange convulsive dance for the next 15 seconds.

and interpreted on NBC, CNN, NYT and other major news outlets.[4] Similar stories, where some user-generated content, channeled by media grabs the attention of the general public are not rare.[5]

Importantly, the emergence of a topic is often helped by specialized sites, also called "aggregators" that focus on identifying potential "hits" from the many topics that appear on the Web. These sites quickly link to sources related to the identified topics, add their own (sometimes controversial) editorial content and then earn advertising revenues from Web traffic as these stories become mainstream. Aggregators do not restrict their targets to the Internet, rather they search for any potential topic that draws exceptional interest from the public, including the news stories of traditional media outlets such as newspapers or television news programs. In fact, the name "aggregator" reflects the fact that many see these sites as, essentially, pirating content from traditional media outlets. The aggregators' success critically depends on two factors. First, on their capacity to rapidly identify popular themes that are likely to become mainstream. To this end they use special technology that monitors early tick-up of Web traffic for newly posted content on various social media or other websites. The second key success factor for aggregators is to be able to generate enough content on their site to become the "go-to-place" for the particular topic in question. Some of this content is created by agile integration or referencing of existing sources but also, by the creation of original content, often by the use of large numbers of freelance 'journalists'. Aggregators are numerous. Beyond Gawker mentioned above, well-known sites include Demand media, a site that uses technology to monitor the public's search behavior to identify emerging themes. BuzzFeed, another popular site identifies "hot content" by monitoring sharing behavior on popular social media sites, such as Facebook or Twitter. The Business Insider is a similar site specialized in providing business news. A number of new entrants (e.g. Flipboard or NowThisNews) exploit the growing demand for

[4] See http://www.cnn.com/video/#/video/international/2013/02/14/new-dance-crazy-harlem-shake.cnn or http://www.nytimes.com/2013/02/20/arts/music/macklemores-thrift-shop-and-baauers-harlem-shake.html?_r=1&adxnnl=1&adxnnlx=1361400282-pk6WpMtpfVQ6n4hnMw3TtQ&.

[5] See, for example, the well-documented case by Deighton and Kornfeld (2010) concerning the broken guitar of a United Airlines passenger.

on-demand video, providing video news to compete with traditional television news channels. Barriers to entry are generally low, which means that competition between these sites is fierce. The emergence of aggregators also represents a huge challenge for traditional news providers in a world where user-generated content competes with the traditional ways to generate worthy news for readers/viewers. In fact, the line between traditional and new media is increasingly blurred as even iconic examples of the respective categories copy each other's strategies. For example, BuzzFeed has recently hired dozens of reporters, many of them previously employed by traditional news outlets. Similarly, traditional media, including the New York Times, has realized the need to link to external sources of content to remain relevant.

Beyond the proliferation of competitors, the dynamics of competition has also changed with the emergence of the Internet. Competing news providers are essentially in a permanent contest to be among the few sites that attract the attention of the public. Time is of essence for publishing the 'scoops' as the topic only attracts a disproportionate number of viewers for a limited amount of time, after which other news sites can catch up. If done well however, then for a short period (typically a day), the site can quickly increase its traffic, which in turn raises advertising revenues. This is quite different from the traditional, subscription-based or brand-based competition between news providers where, due to switching costs, the publisher (e.g. a newspaper) is somewhat shielded from occasionally missing out on a story. Moreover, in the traditional model, substantial fixed costs are needed for the generation of content (e.g. the cost of maintaining a crew of reporters), which represent barriers to entry leading to a much smaller set of competing firms. In the online news market entry barriers are relatively low as Internet users, rather than reporters generate content. Moreover, consumer switching costs are low and viewers can quickly converge to the sites with the most relevant 'topics'.[6] Traditional news providers have taken notice and they have also launched their own 'Web corners' on their sites

[6]Our model describes some aspect of competition among traditional media firms who had to pick topics from a relatively large set generated by upstream information providers such as the Associated Press or Reuters. In this respect, we argue that the "contest nature" of competition has significantly increased for the news industry, as did the number of competing outlets.

(see for example CNN's 'distraction videos'), which try to rely on similar forecasting techniques to identify emerging popular topics.

The central question of this paper is how these competitive dynamics are likely to influence the public agenda, that is, what type of news will emerge from the massive amount of content available? In particular, will competition focus firms on a relatively few important topics or will such competition lead to the vast proliferation of published topics with a fragmented news scene? How will the distribution of content and the number of competitors influence the nature of published news and its diversity? How will asymmetries across news providers influence these outcomes? To answer these questions, we develop a generalized contest model in which firms have to choose one news item from a large set of items with varying prior probabilities of success. We allow for the simultaneous and/or correlated success of multiple items. Importantly, we model *all* relevant ways in which 'winning sites' (in the sense of reporting on eventually popular topics) may share the reward for successful news. We also explore the effect of differences across firms, first by assuming that some firms have a loyal customer base (e.g. a brand) and, next, by endowing a subset of firms with a better capability for forecasting the news items that may become eventually popular.

We generally find that the variety/diversity of topics as well as the weight given to 'marginal' (*a priori* unlikely) topics increases the more the topics' prior follows a "fat-tail" distribution and the more correlated the success of *a priori* likely topics are. More importantly, we find a non-trivial effect for the number of competing firms. Interestingly, as long as the contest is "not too strong" among sites, their choice of published topics is concentrated on the news with the highest prior success probabilities and this is even more so the more sites enter the market. In other words, in this case, increased entry actually reinforces the concentration of news. In contrast, when the intensity of the contest among sites is beyond a certain threshold, competition tends to rapidly increase the fragmentation of published news: as the number of competing publishers increases, more and more *a priori* unlikely topics are reported resulting in a large diversity of published topics. Next, we focus on differences across firms (i) in their capability to identify potential 'hits', as well as, (ii) in the size of their loyal customer base (brand value). We find that when some firms have better technology to forecast the popularity of topics, then, surprisingly, overall diversity of news published by the remaining firms declines as these firms tend to take refuge in publishing on 'safer' topics. In contrast, when a subset of firms have extra revenue from a published 'hit' from loyal users then these 'branded' publishers tend to be conservative in their choice of topics as their loyal customer base represents 'insurance' against the contest. In contrast, the diversity of news published by unbranded outlets increases as unbranded publishers tend to avoid branded ones by putting more weight on *a priori* unlikely stories. These results are consistent with anecdotal evidence in the news industry and they also conform the broadly observed evolution of diversity in the public agenda (see, e.g. McCombs and Zhu (1995)). In a final analysis, we consider endogenous success probabilities. It is widely accepted that the media often 'makes the news' in the sense that a topic may become relevant simply because it got published. Interestingly, such a dynamic has an ambiguous effect on the diversity of published topics. If the contest is very strong then it results in a concentrated set of *a priori* likely topics. When the contest is moderate then the diversity of topics may be higher depending on the number of competing outlets.

The paper is organized as follows. In the next section, we summarize the relevant literature. This is followed by the description of the basic model and its analysis where we present the main results. Next, we explore the impact of asymmetries across firms and endogenous success probabilities. The paper ends with a discussion of the results and concluding remarks. To facilitate reading, all proofs are relegated to the Appendix.

2 Relevant literature

The topic of this paper is generally related to the literature on agenda setting (see McCombs (2004) for an excellent recent review) that studies the role of media in focusing the public on certain topics instead of others. It is broadly believed that agenda setting has a greater influence on the public than published opinion whose explicit purpose is to influence the readers' perspective. As the famous saying by Bernard Cohen (1963) goes: "The media may not be successful in telling people what to think but they are stunningly successful in telling their audiences what to think about". The literature examines the mechanisms that lead to the emergence of topics and the diversity of topics across media outlets. In particular, McCombs and Zhu (1995) show that the general diversity of topics as well as their volatility has been steadily increasing over time. The general focus of our paper is similar: we show that the nature of competition is an important mechanism affecting the emerging public agenda in the news.

Agenda setting is also addressed in the literature studying the political economy of mass media (see Prat and Stromberg (2013) for an excellent review).[7] The standard theory states that media coverage is higher for topics that are of interest for (a) larger groups, (b) with larger advertising potential, and (c) when the topic is journalistically more "newsworthy" and (d) cheaper to distribute. While there is little empirical evidence to support (b), the other hypotheses are generally supported (see Stromberg (2004) and Snyder and Stromberg (2010), among others). Hypotheses (c) is particularly interesting from our standpoint. Eisensee and Stromberg (2007) show that the demand for topics can vary substantially over time. For example, sensational topics of general interest (e.g. the Olympic Games) may crowd out other 'important' topics (e.g. natural disasters) that would be covered otherwise. This supports the general notion that media needs to constantly forecast the likely success of topics and select among them accordingly. Our main interest is different from this literature's as we primarily focus on media competition as opposed to what causes variations in demand. Taking the demand as given, our goal is to understand how the competitive dynamics between media firms distorts the selection of topics, which then has a major impact on agenda setting.

[7]While not our focus, this literature also addresses the related issue of media bias and how it is affected by media competition – see Mullainathan and Shleifer (2005) and Xiang and Sarvary (2007) for relevant analytical models and Gentzkow and Shapiro (2010) and Larcinese and Snyder (2011) for empirical evidence. The latter paper is interesting because it shows how newspapers can achieve bias by overrepresenting favourable topics to politicians close to their voters.

As such, the paper also relates to the growing literature on media competition where the strategic variable is content quality broadly defined. The primary focus of empirical research is on how media concentration affects the diversity of news both in terms of the issues discussed in the media as well as the diversity of opinion on a particular issue. For example, George and Oberholzer-Gee (2011) show that in local broadcast news, "issue diversity" grows with increased competition (as measured by the number of local TV stations) even though political diversity tends to decrease. Franceschelli (2011) studies the impact of the Internet on news coverage, in particular the recent decrease in the lead-time for catching up with a missed breaking news. He argues that missing the breaking news has less impact, as the news outlet can catch up with rivals in less time. This might lead to a free-riding effect among media outlets, where there is less incentive to identify the breaking news. Both of these papers have consistent empirical findings with our results/assumptions. On the theory front, a recent paper by Dellarocas et al. (2013) examines the competition between media outlets that can link to each other and the role of news aggregators. They show that the presence of links and aggregators can substantially alter the nature of competition and lead to better or worse quality content.[8] Yildirim et al. (2013) examine the decision of news media to include user-generated content in their online edition. Although such an addition reduces profits of competing news providers and increases the bias online, it mitigates the bias in the print edition. Finally, Xiang and Soberman (2013) model the competitive design of news when the demand for topics is uncertain. However, their focus is on editorial design, which allows consumers to process more information. They show that a better design is not always optimal unless the news program is "complex" in the sense that it contains a large number of topics.

In terms of the analytical model, we rely on the literature studying competitive contests among forecasters. For example, Ottaviani and Sorensen (2006) use a similar framework to model competition among financial analysts. Our model is different in that we explore in more

[8]Mayzlin and Yoganarasimhan (2012), consider a somewhat similar problem with competing blogs that build reputation by identifying "scoops" for their readers. While the general issue is similar to ours, their focus is on exploring how firms rely on content sharing to mitigate missing out on relevant news.

detail the structure of the state space, we generalize the contest model and extend it in a variety of ways, most notably by analyzing asymmetries across firms.

3 Base model

We model the competition between news providers as a contest, where each outlet tries to identify and publish the topics that will become the central interest for the public. We assume that these topics will attract disproportionate attention from viewers corresponding to extraordinary advertising revenues, which essentially constitutes the prize of the contest that is shared across the 'winning' sites, those who report on this topic. The challenge for the media outlets is that, in addition to accurately forecasting the likely success of stories from a vast amount of content, they also need to try to identify unique stories that other sites did *not* publish. The contest nature of competition comes from the fact that successful sites need to share the audiences if they all identified the same story(ies).

To formalize this setup assume N competing news providers and K topics with p_k probabilities ($k = 1,\ 2, \dots,\ K$), where p_k measures the prior probability that topic k will become successful news (i.e. capture the attention of the public). One could imagine, for example, that the K topics are pieces of content appearing on the Internet. We assume that p_k and the entire joint distribution of the K events is exogenous and common knowledge across all outlets, that is, we assume that media firms have identical technology to forecast the likely success of available content.[9] To denote the joint distribution of topics, we use P_S for the probability that exactly events in set S become successful for a $S \subseteq \{1, 2, \dots, K\}$.

Without loss of generality, we rank events in decreasing order of prior probabilities ($p_{k_1} > p_{k_2}$ if $k_1 < k_2$).[10] Note that we allow for $\sum_k p_k \geq 1$, i.e. it is possible that several topics may become mainstream news simultaneously. We assume that the media outlets can only choose one of the K topics for publication. This reflects the idea that the news provider can only select a small

[9] In §5.2, we relax this assumption and explore asymmetry in sites' forecasting capabilities.

[10] For technical convenience, we assume that all p_k probabilities are different, but this is not a crucial assumption.

set from the large pool of topics for publication and needs to put quite some effort in becoming a "relevant" destination for these topics.

As an illustration, consider the simplified problem, where $K = 2$, i.e. the set of topics is $\{1, 2\}$ and, without loss of generality, $p_1 > p_2$. To complete the problem, assume that the probability that both events become mainstream is $P = P_{\{1,2\}}$. Then, the probability that only event 1 is successful is $p_1 - P$ and the probability that only event 2 is successful is $p_2 - P$. The probability that no event is successful in capturing the interest of the public is then $1-p_1-p_2+P$. One way to think about this simplification is that there are a few "important" topics that have a high probability of becoming mainstream (major wars, sport events, elections, etc.) Besides these however, there is a very large number of "marginal" topics (say, the long tail of Web content) that have a low probability to become mainstream. A central question is: how likely is it that a firm chooses from the low-probability topics? In this context, $p_1/p_2 > 1$ measures how skewed the prior distribution of content is, i.e. p_2 represents the mass corresponding to the long tail of this content. In the general case of $K > 2$, we explicitly model the individual long tail topics.

To model the reward of providers that choose a topic that becomes a success, let us first normalize the reward that a single publisher gets when picking an event that becomes the only successful topic to 1. This is the total value of the audience of a single successful topic. We introduce two types of parameters to capture competition between topics and publishers. Let γ_ℓ measure the value of the audience of one out of ℓ successful topics capturing the competition between topics, where $\gamma_1 = 1$ and γ_ℓ is (weakly) decreasing in ℓ. Furthermore, let β_n measure the competition between publishers for this audience. When one out of n publishers picks a successful topic, it receives β_n portion of the total reward for that topic. We assume that $\beta_1 = 1$ and $\beta_n > 0$ is (weakly) decreasing in n, converging to $\beta_\infty = \inf\{\beta_n, n \geq 1\}$.[11]

Formally, let y_i denote site i's choice and let n_k denote the number of sites that choose k.

[11] We assume that β_n is positive for technical convenience. The results are similar when we allow $\beta_n = 0$.

If y is any one of the ℓ topics that became mainstream then site i's payoff is

$$\pi_i = \begin{cases} \beta_{n_{y_i}} \gamma_\ell & \text{if } y = y_i \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

For example, in a pure contest where $\beta_n = 1/n$, a fixed reward (γ_ℓ) is divided equally between the sites who publish a single topic that is among the ℓ successful ones. If multiple topics become successful, then the total reward accrued to the sites that publish one of these is smaller and more so the more topics become successful, i.e. more successful news items reduce the demand for any individual one. It is possible however, that the reward accrued to all successful sites is larger than 1 (that is, when $\ell\gamma_\ell \geq 1$). In other words, a larger number of interesting news events may increase total media demand. The β_n sequence measures how competition between publishers increases with more players. A constant $\beta \equiv 1$ implies no competition between sites, whereas a sharply decreasing β_n series describes a market that gets very competitive with more players.

4 Analysis

For a single news provider ($N = 1$) the problem is trivial: choose $y = 1$, the topic with the highest prior probability. When there are multiple providers, however, there might be an incentive to choose from the topics with the lower priors because if providers choose different topics, the prize needs to be shared with fewer competitors. Moreover, the lower the prior, the fewer the sites willing to publish the corresponding event. One can show that there are many pure-strategy equilibria in which different sub-groups coordinate on different topics with group sizes being larger for higher probability topics. Given the multiplicity of pure-strategy equilibria and the difficulty of coordination across a large number of firms, the relevant equilibrium is in mixed-strategies. Indeed, a symmetric mixed-strategy equilibrium allows us to better characterize the expected distribution of topics resulting from firms' choices. Let $q_k^{(N)}$ denote the equilibrium probability that a firm chooses topic k when there are N players. In order to present our main results, let us first define

$$v_k = p_k - \sum_{\{k \in S,\ |S| \geq 2\}} (1 - \gamma_{|S|}) P_S \tag{2}$$

and let $O(.)$ be a decreasing ordering by v_k such that $v_{O(1)} \geq v_{O(2)} \geq \ldots \geq v_{O(K)}$. As we show below, v_k represents the value of a topic and players choose between topics according to a decreasing order of value:

Proposition 1 *The game with $N \geq 1$ players has a unique symmetric equilibrium in mixed strategies with the following properties.*

1. *If $v_i > v_j$ holds for a pair of topics (i, j), then $q_i^{(N)} > q_j^{(N)}$ for any $N \geq 1$.*
2. *There exists an increasing $\overline{K_N}$ series such that topic $O(j)$ is chosen with positive probability, i.e. $q_{O(j)}^{(N)} > 0$, if and only if $j \leq \overline{K_N}$.*
3. *If topic k is chosen with positive probability, then $v_k > v_{O(1)}\beta_N$.*

The proposition has two key messages. First, it shows that the "value" of topics for firms, measured by v_i does not necessarily correspond to their prior probabilities of success. Rather, it is also a function of their correlation structure, i.e. the likelihood that some of them become successful *together*. We get $v_i = p_i$ only if $\gamma_\ell = 1\ \forall \ell$ (i.e. multiple topics becoming successful simply multiplies the overall demand for news) or if $P_S = 0\ \forall |S| \geq 2$, (i.e. topics are mutually exclusive). While a reversal of value ($v_i > v_j$ when $p_i < p_j$) cannot happen for $K = 2$, it is possible for $K \geq 3$. As an example, assume three topics, i, j and k with $p_i > p_j > p_k$, where two of the events are highly correlated: say, if topic i becomes successful it is likely that topic j becomes successful too ($P_{\{i,j\}}$ is large). If event k is negatively correlated with the other two events and γ_2 is not too large, then topic k may attract disproportionate share of choice from firms even with a low prior. This happens because event k is likely to happen alone and, therefore, the demand will not be divided between two events. In this case, the probability that firms choose k in equilibrium may exceed k's prior probability, p_k. This simple example shows that under competition, the correlation structure between topics may also have a major role in determining what gets published from a large number of topics. In particular, in a contest, even quite unlikely topics may make it to the news if they are "unique" compared to others, in the sense that they are likely to become successful independently from other topics.

As an example, think of general themes like a presidential election or the Olympic games. Most 'relevant' news related to these themes (e.g. the state of the economy, political news, gold medals won), are likely to attract the attention of the public, i.e. their joint probability of "success" is high. Our model predicts that with increased competition between news outlets 'irrelevant' news (e.g. gossip about the presidential candidates or the missteps of athletes) may get over-represented in the news.[12] Often news consists of signals or forecasts about future events that are of general relevance to the public. For example, news may report poll results forecasting the outcome of elections or an important vote in congress, etc. If the success of these news is related to the likelihood of them forecasting the truth then our model may explain why "controversial" or "surprising" forecasts about future events may be over-represented in the media: a strong prior about the future event would make most forecasts (those consistent with the prior) correlated (i.e. to be true *together*). In contrast, while the success of an inconsistent poll is low, its success is independent from those forecasts. For instance, a poll predicting the failure of the front-runner in the presidential election is likely to be wrong and, therefore, less likely to become a successful story but it is also likely that, *if successful* it is a unique story. This latter aspect provides an extra incentive for competing media firms to report it.

The second insight from Proposition 1 is related to the level of competition, measured by the number of news outlets, N in the model. The more firms compete in the contest for successful news the more there is a chance for "low value" (either unlikely or highly correlated) topics to make it to the news. As $N \to \infty$, all K topics will be published with positive probability

[12] The last presidential election in France is a good example. Most news outlets covered totally irrelevant topics (a notable one was the debate over the correct labeling of halal food in boucheries) in a context where the country was facing a major economic crisis. *The Economist* has devoted its front page to the election with the title: "France in denial".

(although we will see below that this probability may be very small and actually further decrease with even larger N). In other words, competition tends to increase topic diversity, the more so the more the contest is competitive between the publishers.

The next proposition explores the evolution of probabilities or the "attention" that topics get in the news as a function of competition, defined by the *nature* of the contest between publishers. To measure the increase in competition with more players we define $r(\beta) = -\log_2\left(\lim_{n\to\infty}\frac{\beta_{2n}}{\beta_n}\right) \geq 0$, measuring how fast the β sequence converges to 0.[13] For example, if $\beta_n = 1/n^s$ then $r(\beta) = s$. When β_n does not converge to 0, then $r(\beta)$ is clearly 0, but even if it decreases slowly, such as when $\beta_n = 1/\log n$, we get that $r(\beta) = 0$. On the other extreme, when β_n decreases exponentially or even faster, $r(\beta) = \infty$. Depending on how fast β_n decreases, we get substantively different results as the number of players approaches infinity.

Proposition 2 *As $N \to \infty$, the equilibrium mixing probabilities converge as follows.*

1. *If $r(\beta) = 0$, then $q^{(N)}_{O(1)} \to 1$ and $q^{(N)}_{O(i)} \to 0$ for any $i \geq 2$.*
2. *If $r = r(\beta) < \infty$, then $q^{(N)}_k \to (v_k)^{1/r}/\sum_{j=1}^{K}(v_j)^{1/r}$ for any $1 \leq k \leq K$ topic.*
3. *If $r(\beta) = \infty$, then $q^{(N)}_k \to 1/K$ for any $1 \leq k \leq K$ topic.*

The results show how the level of competition between publishers affects topic choice and coverage. We have seen that a higher number of players always leads to more topics covered with positive probability, but the evolution of topic choices depends on how competitive the contest is. When there is absolutely no competition, i.e., $\beta_n \equiv 1$, players always chose the most promising topic regardless of the number of players. More interestingly, even if the contest is only mildly competitive (i.e. if competition does not increase very much with the number of players: $\lim_{n\to\infty}\beta_n > 0$ or even if $\lim_{n\to\infty}\beta_n = 0$, but slowly), in the limit everyone will choose the topic with the highest prior probability.[14] When competition is more intense, e.g.

[13] We assume that β is such that $\lim_{n\to\infty}\frac{\beta_{nx}}{\beta_n}$ exists for any rational $x > 1$.

[14] This outcome does not contradict the result of Proposition 1. When $r(\beta) = 0$ the evolution of probabilities follows an intriguing pattern. As N increases additional topics with lower and lower v_i-s obtain a positive weight

$\beta_n = 1/n^r$, the limit will be a diverse set of choices with probabilities proportional to $v_k^{(1/r)}$. To better illustrate this case, let us explore the simplest case with $\beta_n = 1/n$ and only two topics ($K = 2$) where $p_1 + p_2 = 1$, $P = P_{\{1,2\}} = 0$, and $p_1 > p_2$, i.e. only one of the two topics becomes mainstream for sure. Then, in the mixed-strategy equilibrium, firms will choose the less likely topic with probability $2 - 3p_1$, which is larger than 0 as long as $p_1 < 2/3$. However, as the number of firms increases, the probability of choosing the less likely topic increases until it reaches p_2 when $N = \infty$. Similarly, the probability of choosing the more likely topic decreases. Figure 1 shows how increased competition reduces the probability of choosing the more likely event as a function of its prior, p_1. Note that for $N = \infty$ this probability is exactly p_1. According to our results, for the more general case with $K = 2$, $\frac{q_2}{q_1} = \frac{p_2-(1-\gamma)P}{p_1-(1-\gamma)P}$. As we have seen before, for $P = 0$ or $\gamma = 1$ this yields that, with many firms, the relative probability of publishing each topic corresponds to the relative proportion of prior probabilities.

The final case in Proposition 2 is also interesting. It considers a very tough contest in which the reward dissipates very quickly as multiple sites choose the same successful topic(s). The Proposition says that, in this case, news outlets completely randomize their choice of topics by putting equal weight on each of them leading to extreme topic diversity in the news. A simple example for such an extreme contest is one where a firm's reward falls to 0 as soon as another firm also reports on the same topic. One way to interpret this situation is to assume that the reward consists of advertising revenues and there is Bertrand competition between news sites on the advertising market. In this case, all topics would be equally represented in the news.

In summary, our model predicts that with a competitive contest, the more firms enter the market the less players choose topics that are likely to succeed *a priori* and topic diversity tends to increase in the published news. Note, however, that even with an infinite number of firms, the average representation of topics may not correspond to the marginal distribution of priors, as correlations distort the representation of topics. Only when the topics are uncorrelated (i.e.

in firms' choice. However, after obtaining a positive weight, this weight decreases, converging to 0 with more firms entering the market. Therefore, from a practical perspective, when the contets is not too competitive, the concentration of topics increases with entry.

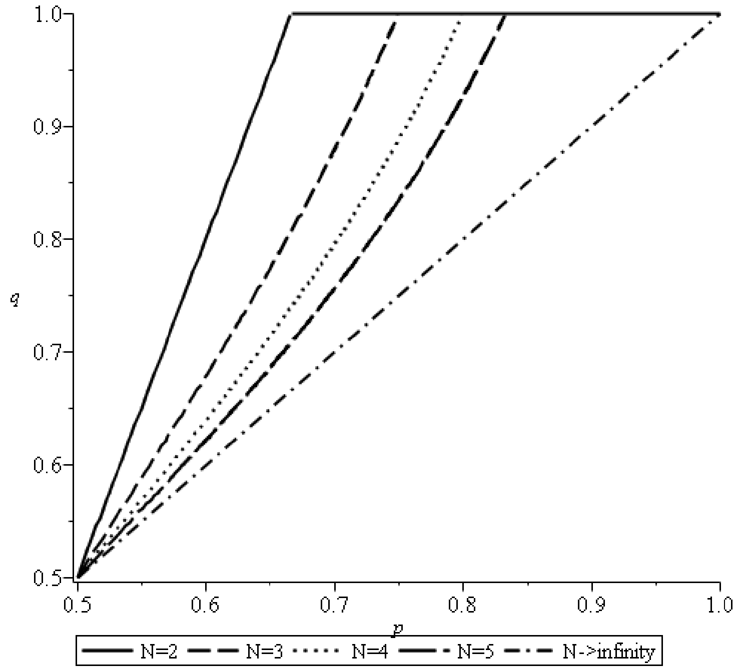

Figure 1: Equilibrium probability of choosing topic 1 for $N = 1, 2, 3, 4, 5$ and $N \to \infty$.

when $v_i = p_i$), do an infinite number of firms replicate the prior distribution of topics. In the final case of the Proposition, i.e. when the contest is extremely competitive, it drives players to differentiate as much as possible and, in the limit, they choose each topic with the same likelihood, irrespective of the topics' priors. Figure 2 compares our results for different values of r, showing the different outcomes depending on how competitive the contest is.

We have seen how the nature of the contest and the number of players affect the choice of topics between different players, generally leading to more diversification and differentiation in case of more players. It is also important to consider how a change in the available topics changes players' choices. Let us assume that the set of available topics is modified in a way that topics $1, 2, 3, \ldots, K-1$ remain the same, but topic K is split into two separate topics such that the value of the split topic is equal to the sum of the values of the two individual topics. A good example is the case of mutually exclusive topics, where the value of a topic is simply the probability of it becoming successful. A topic can be refined into two versions such that

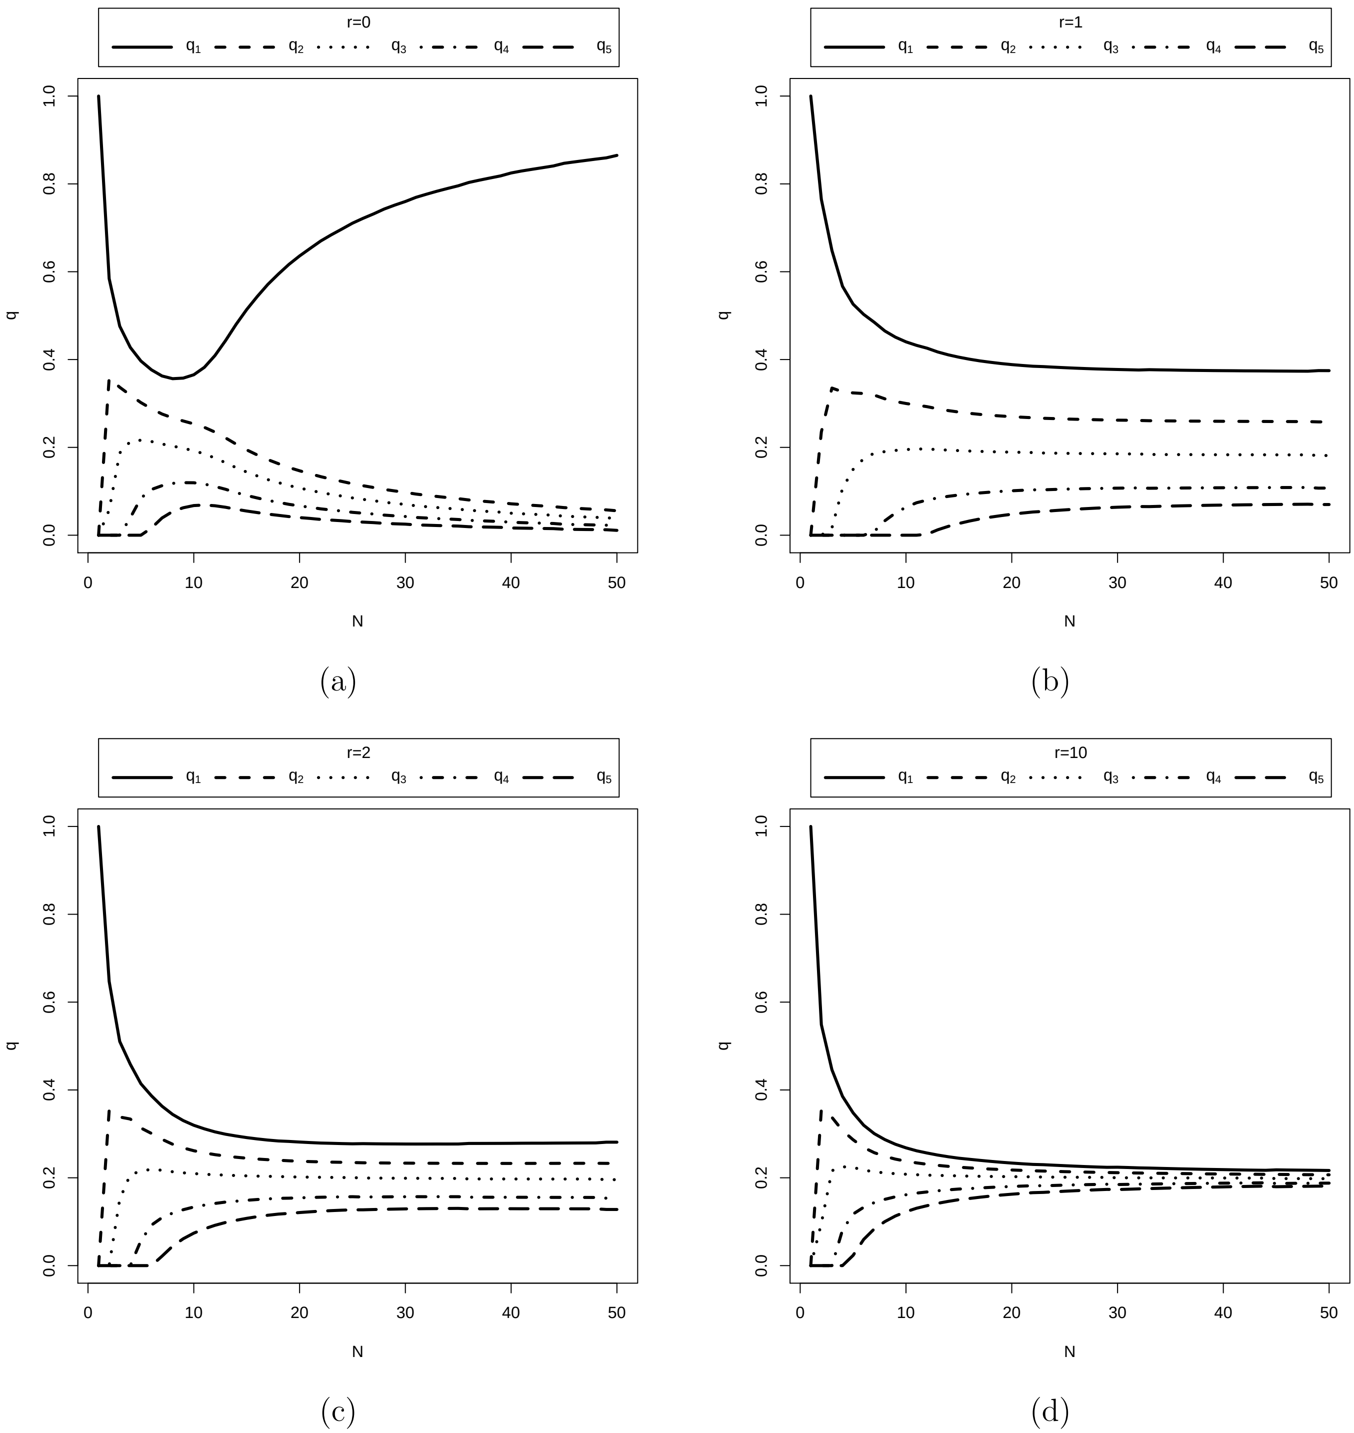

Figure 2: Mixing probabilities for $K = 5$ with $v_1 = 0.5, v_2 = 0.35, v_3 = 0.25, v_4 = 0.15, v_5 = 0.1$. As N increase, more topics are choosen with positive probability. However, when $r = 0$ all probabilities except q_1 converge to 0. When $r > 0$, the mixing probability for large N's is proportional to $v_k^{(1/r)}$, eventually leading to equal probabilities as $r \to \infty$.

players can pick one version or the other. As the following corollary shows, splitting topics in such a way decreases the likelihood that players pick them.

Corollary 1 *Assume $\beta_n = 1/n$ and let $q_i^{(N)}$ denote the symmetric equilibrium probabilities of choosing between the K topics with values $v_1, v_2, ..., v_K + v_{K+1}$ and let $s_i^{(N)}$ denote the equilibrium probabilities of choosing between the $K+1$ topics with values $v_1, v_2, ..., v_K, v_{K+1}$. Then $s_K^{(N)} + s_{K+1}^{(N)} < q_K^{(N)}$ for any $N \geq 1$.*

5 The role of asymmetries

In the basic model, we studied a general setup with symmetric publishers. In this section, we explore the role of asymmetries. First, we consider how branded publishers with a loyal consumer segment behave and how this impacts the rest of the players. Second, we model the potential differences in players' abilities in determining which topic is going to be successful.

5.1 Branded providers

To account for branded news providers, we assume that brand value manifests itself in a loyal segment that only consumes a given provider's stories. Each branded provider has a loyal segment of size $\lambda > 0$, and each of these consumers provides a revenue of 1 if the branded provider's story is successful and 0 otherwise. We modify the payoff function in (1) for branded publishers to

$$\pi_i^B = \begin{cases} \beta_{n_{y_i}} \gamma_\ell + \lambda & \text{if } y = y_i \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Recall that y_i above denotes the choice of player i, whereas y is one of ℓ successful topics. Therefore, the payoff includes the same competitive component as before from non-loyal consumers, but also includes a unit profit from all loyal customers. The payoff of all non-branded sites is identical to π_i as before in (1). For simplicity we assume that all branded sites have the same amount of loyal consumers and we use α to denote the proportion of branded providers. It is not hard to map this setup to today's situation in the news media: branded news providers are represented by traditional firms (e.g. the New York Times, Washington Post, etc. for newspapers or CNN, NBC or Fox News for television broadcasters) while entering, mostly online news outlets (BuzzFeed, Huffington Post) represent unbranded providers.

To compare this setting to our main findings, we search for equilibria in mixed strategies where the strategies only depend on the type of provider and do not differ within the set of branded players or within the set of non-branded players. We call this a symmetric equilibrium. For a first look, let us consider the case with only one branded player, say player 1. It is clear that if non-branded players are indifferent between two topics k_1 and k_2 then the branded player will choose the topic with the higher prior probability, that is k_1 iff $p_{k_1} > p_{k_2}$. As such, the branded player will chose topic 1, the topic with the highest prior probability. This, in turn will incentivize non-branded player to avoid topic 1. We state the result in a general fashion below.

Proposition 3 *There exists $\overline{\alpha} > 0$, such that if $0 < \alpha < \overline{\alpha}$, then the game has a unique symmetric equilibrium, where all the branded players always choose topic 1, that is $q_1^B = 1$. The non-branded players' q_k^{NB} mixing probabilities satisfy $q_1^{NB} < q_1$ and $q_k^{NB} > q_k$, for any $k \geq 2$, where q_k is the equilibrium with only non-branded players as in Proposition 1.*

The result shows that when the numerosity of branded players is not too high, they all bet on the topic with the highest prior probability so that they do not disappoint their loyal customers. The concentration of branded players leads non-branded providers to avoid this topic and pick others.

We can obtain a more precise picture of the magnitudes by examining the case of $N \to \infty$.

Corollary 2 *Assume $r = r(\beta) > 0$ and $\alpha < (v_1)^{1/r} / \sum_{j=1}^{K} (v_j)^{1/r}$. Then, as $N \to \infty$*

$$q_1^{NB(N)} \to \frac{(v_k)^{1/r} / \sum_{j=1}^{K} (v_j)^{1/r} - \alpha}{1 - \alpha},$$

$$q_k^{NB(N)} \to \frac{(v_k)^{1/r} / \sum_{j=1}^{K} (v_j)^{1/r}}{1 - \alpha},$$

for any $k \geq 2$.

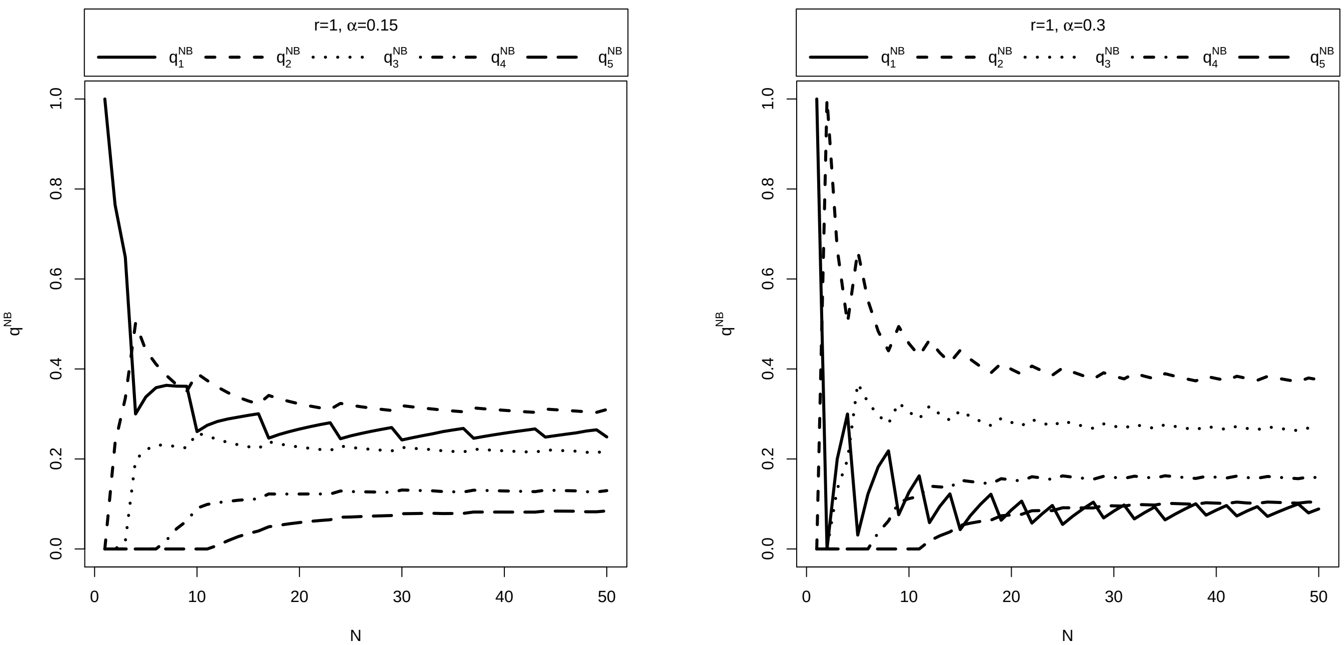

(a) (b)

Figure 3: Non-branded provider's mixing probabilities for $K = 5$ with $v_1 = 0.5, v_2 = 0.35, v_3 = 0.25, v_4 = 0.15, v_5 = 0.1$ for different values of α (0.15, 0.3). A high proportion of branded providers makes non-branded providers turn away from high value topics.

The corollary reveals that the α proportion of branded providers dominate topic 1, thereby increasing the likelihood of other topics chosen by the non-branded providers. The more branded providers there are the more non-branded ones turn away from topic 1 as shown by the mixing probabilities that are increasing in α for topics $k \geq 2$. Figure 3 illustrates the results for $\alpha = 0.15$ and $\alpha = 0.3$.

Note that although topic 1 has the highest prior probability it is not necessarily the highest value topic in the absence of loyal customers. Depending on the correlation structure these two topics can differ. When they do, we observe more of a horizontal differentiation with branded providers choosing the most likely topics, whereas non-branded providers choosing the most likely *but also* unique topic. If the two coincide then the differentiation is more vertical since non-branded providers choose topics that are less likely to be successful.

The result described by Proposition 3 and Corollary 2 is consistent with the casual observation that traditional (branded) media is relatively conservative in their choices of top stories. While recently, even established media firms ventured into publishing less traditional news on their websites, these are typically relegated to special sections that are clearly suggested to be taken "more lightly" (see CNN's "Distraction videos" mentioned earlier). In contrast, new media sites are much more venturesome in their editorial process often reporting stories that could be easily qualified as 'rumor' or 'gossip'.

5.2 Predictive ability

So far, we assumed symmetric information across publishers. However, some publishers may have an advantage in determining which topic would become successful in the future. To study information asymmetry, we start from our basic model using the payoffs given in (1) but we assume that a μ proportion of players can perfectly predict which topic(s) will be successful. The remaining $1-\mu$ proportion has the same prior information as before. For tractability, we assume that the topics are mutually exclusive,[15] that is, $P_S = 0$ for any $|S| \geq 2$. Again, this setup can be easily mapped to today's changing news media landscape. Here, media firms/sites with better forecasting ability clearly correspond to new online entrants (e.g. the Huffington Post or Buzzfeed), who are often very open about their superior technology in terms of predicting the success of emerging stories using social media or search engines.

Naturally, in this setting, high-ability publishers will choose a topic that will eventually become successful as they can perfectly predict success. The question is how this behavior affects the topic choice of the remaining $(1-\mu)N$ players.

Proposition 4 *The game has a unique symmetric equilibrium in mixed strategies for less-informed players. For any μ, there exists a $\overline{K_N}(\mu)$ series, increasing in N, such that topic $O(j)$ is chosen with positive probability by less-informed players, i.e. $q_{O(j)}^{L(N)} > 0$, if and only if $j \leq \overline{K_N}(\mu)$. If topic k is chosen with positive probability, $v_k > v_{O(1)}\beta_N/\beta_{\mu N+1}$.*

[15] This implies $v_k = p_k$, but for notational consistency, we still use the v_k values.

The special case of $\beta_n = 1/n$ illustrates our results best. When $\mu = 0$, that is, all players have the same forecasting ability, the necessary condition for a topic to be chosen is that topic k's value, $v_k > v_{O(1)}/N$. However, as a positive proportion of μ players have high abilities, the condition becomes $v_k > v_{O(1)}(\mu N + 1)/N$, essentially ruling out topics with a value less than μ times the highest value topics. In other words, the higher the proportion of high-ability players, the less uninformed players chose unlikely topics. As $N \to \infty$, we can derive the mixing probabilities for less-informed players.

Corollary 3 *Assume $r = r(\beta) > 0$. Then, $\overline{K_N}(\mu) \to \overline{K_\infty}(\mu)$, where $\overline{K_\infty}(\mu)$ is decreasing in μ and is defined as the largest integer satisfying*

$$\frac{(v_{O(\overline{K_\infty}(\mu))})^{1/r}}{\sum_{j=1}^{\overline{K_\infty}(\mu)}(v_O(j))^{1/r}} \geq \frac{\mu}{(\overline{K_\infty}(\mu) - 1)\mu + 1}.$$

Furthermore, as $N \to \infty$, for any $j \leq \overline{K_\infty}(\mu)$

$$q_{O(j)}^{L(N)} \to \frac{\left((v_{O(j)})^{1/r} / \sum_{j=1}^{\overline{K_\infty}(\mu)}(v_j)^{1/r}\right)\left((\overline{K_\infty}(\mu) - 1)\mu + 1\right) - \mu}{1 - \mu}.$$

For $j > \overline{K_\infty}(\mu)$, we have $q_{O(j)}^{L(N)} \equiv 0$.

Using the $\beta_n = 1/n$ case again, we can see that with a positive μ, not all topics are chosen by low-ability news providers even if there are a large number of them. When $N \to \infty$, topics that have a relatively low value (in proportion to the other topics) are not in the mix. The threshold is decreasing with μ, that is, the higher the proportion of high-ability players, the fewer, and higher value topics will be chosen by low-ability players, as illustrated in Figure 4. This outcome is, again, consistent with what we observe in practice. News outlets with better forecasting ability are often the ones that report on an unexpected topic (e.g. the 'Gangnam style' phenomenon) as they are more likely to foresee general interest for such topics by the public. In contrast, and maybe as a response, traditional media is much more conservative in its editorial choice of topics.

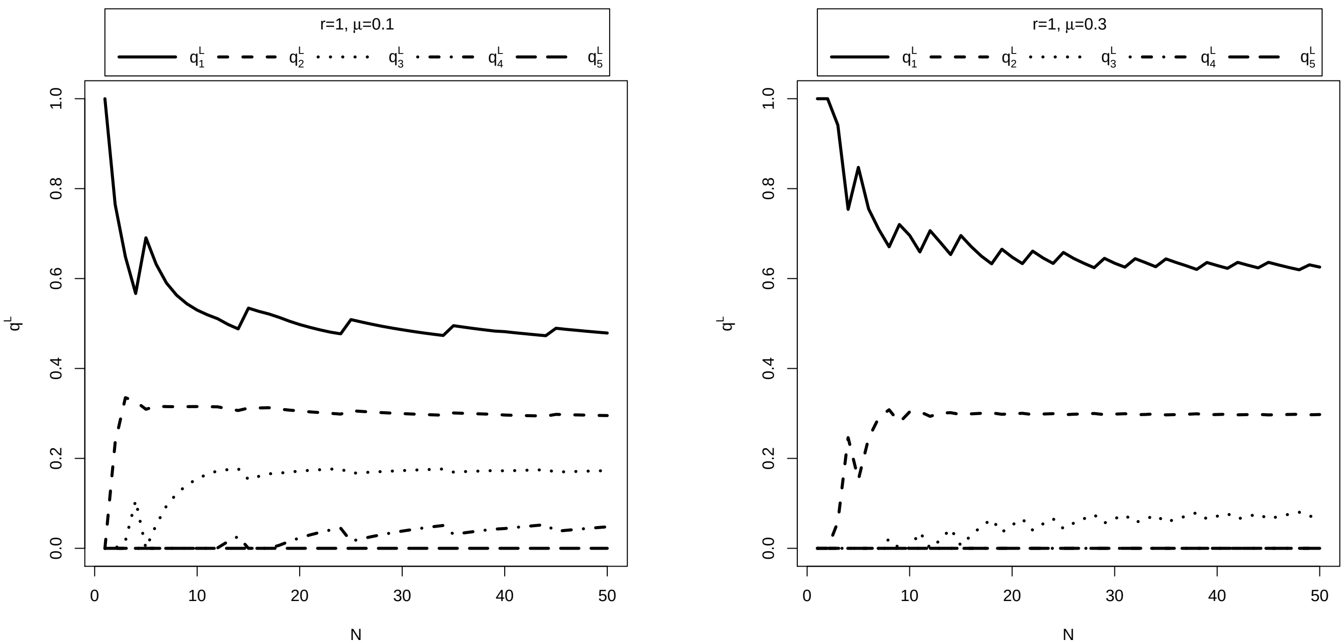

(a) (b)

Figure 4: Low-ability provider's mixing probabilities for $K = 5$ with $v_1 = 0.5, v_2 = 0.35, v_3 = 0.25, v_4 = 0.15, v_5 = 0.1$ for different values of $\mu(0.15, 0.3)$. A high proportion of high-ability providers makes low-ability providers turn away from low value topics.

It is also interesting to ask: what happens if branded providers also have better predictive ability? Which effect dominates in this case? We find that in such a model, brand is dominated by predictive ability. If every branded news provider also has superior predictive ability then brand doesn't matter anymore: every such provider will choose the news that will become successful and all other firms become more conservative in their reporting. There is a caviet however. In our present model, superior predictive ability means that the corresponding firms know for sure which topics will become successful. Clearly, the interplay of these two effects should depend on their relative effectiveness.

6 Endogenous success

The success of a news story often depends, not only on the intrinsic qualities of the story, but also on how much it is reported by providers. Here, we consider the endogenous component of success probabilities that can depend on who chooses to report which story. Recall that in our basic model P_S denoted the probability that exactly stories in set S become successful. We call these purely exogenous success probabilities. At the other end of the spectrum, we define purely endogenous success probabilities, where exactly one topic can become successful and the probability is proportional to the reporting. Let $\tilde{p_k} = n_k / \sum_{j=1}^{K} n_j$ denote the probability that topic k becomes successful in the purely endogenous setting.[16] Then the probability that exactly topics in set S become successful is

$$\tilde{P_S} = \begin{cases} \tilde{p_k} = n_k / \sum_{j=1}^{K} n_j & \text{when} \quad S = \{k\} \\ 0 & \text{when} \quad |S| \geq 2 \end{cases} \tag{4}$$

Most of the time success probabilities are not completely exogenous or endogenous, hence we use a λ parameter to measure the strength of the endogenous component. That is,

$$P'_S := \Pr(\text{exactly topics in set } S \text{ become a success}) = \lambda \tilde{P_S} + (1-\lambda) P_S.$$

[16]The proportional specification is somewhat arbitrary, but we study the effects in relation to the β_n series which is very general. An alternative would be to pick a specific β_n series and a general function for the success probabilities.

Clearly, the case of $\lambda = 0$ corresponds to our basic model, where success probabilities are not affected by players' choices. On the other extreme, when $\lambda = 1$, success is purely determined by which topics players choose. If all of them choose the same topic, that topic becomes a success for certain. Indeed, symmetric pure-strategy equilibria may exist in this case. When all players choose topic k, each of them has a payoff of β_N, since N players have to share the benefits. By deviating to another topic one player could get all the benefits, but only if that topic becomes a success, which happens with probability $1/N$. Therefore, there exists a symmetric equilibrium with all players choosing the same (any) topic if and only if $\beta_N \geq 1/N$. We call this a self-fulfilling equilibrium, since a topic becomes successful only because all players choose it. In a purely endogenous setting these equilibria exists if competition is not very intense relative to how much the players' actions determine success. For tractability, we assume that $\beta_n = 1/n^r$ and look at pure-strategy symmetric equilibria throughout this section. When $\lambda = 1$, self-fulfilling equilibria exist for any topic if and only if $r \leq 1$. In the general case, we first show that such an equilibrium can only exist for the highest value topics.

Lemma 1 *Let $\bar{j}(\lambda, N)$ denote the largest integer such that a pure-strategy symmetric equilibrium where all players choose topic $O(\bar{j}(\lambda, N))$ exists. Then there also exists an equilibrium where all players choose topic $O(j)$ for any $j < \bar{j}(\lambda, N)$.*

The lemma shows that a topic can only be the pure-strategy equilibrium if all topics that are of higher value are also pure-strategy equilibria. We have established above that $\bar{j}(1, N) = K$ if $r \leq 1$ and $\bar{j}(1, N) = 0$ if $r > 1$. For intermediate values of λ we get the following.

Proposition 5 *Assume $0 \leq \lambda < 1$. The threshold $\bar{j}(\lambda, N)$ is increasing in λ.*

1. *If $r \geq 1$, then $\bar{j}(\lambda, N) \leq 1$.*

2. *If $r < 1$, then $\bar{j}(\lambda, N) > 1$ if λ is high enough.*

3. *For each j and high enough λ, there exist $1 < \underline{N}(j) \leq \overline{N}(j)$ such that $j \leq \bar{j}(\lambda, N)$ iff $\underline{N}(j) \leq N \leq \overline{N}(j)$.*

The results describe how and when the coordination on an *a priori* not necessarily high value topic is an equilibrium. First, when competition is intense and $r \geq 1$, only the highest value topic can be chosen by all players and true self-fulfilling equilibria do not exist. Note that when $r = 1$ this is in contrast to the benchmark case of $\lambda = 1$. When competition is less intense and $r < 1$, self-fulfilling equilibria exist when the endogenous component of the success probability is strong enough. Finally, for any given topic (that is not the highest value), an equilibrium with all players chosing that topic exists as long as λ is high enough and N is in an intermediate range. The non-monotonicity in N is a result of the combination of two forces. On the one hand, more players make it more worthwile to go with the mainstream as deviating to an alternative topic is hard when success is strongly determined by the amount of reporting. On the other hand, when there are too many players, even though the probability of reporting the right topic is high for everyone who coordinates, the contest makes it less appealing to participate in this equilibrium.

7 Discussion and concluding remarks

This paper studies competition among news providers who compete in a contest to publish on a relatively small number of topics from a large set when these topics' prior success probabilities differ and when their success may be correlated. We show that the competitive dynamic generated by a strong enough contest causes firms to publish 'isolated' topics with relatively small prior success probabilities. The stronger this competitive dynamic (either because of a larger number of competitors or because the contest forces firms to share a larger proportion of the winners' reward) the more diverse the published news is likely to be. Applied to the context of today's news markets characterized by increased competition between firms, new entrants and reduced customer loyalty, we expect a more diverse set of topics covered by the news industry. While direct evidence is scarce, there seems to be strong empirical support for the general notion that the public agenda has become more diverse over time while also exhibiting more volatility McCombs and Zhu (1995). This general finding is consistent with our results. While diversity of news may generally be considered a good thing, agenda setting, i.e. focusing the public on a few, worthy topics (arguably a core function of the news industry) maybe impaired by increased competition.

In a next step, we explore differences across news providers and find that branded outlets with a loyal customer base are likely to be conservative with their choice of reporting in the sense that they report news that is *a priori* agreed to be important. Facing new competitors with better forecasting ability also makes traditional media more conservative. In sum, if the public considers traditional media and not the new entrants as the key players in agenda setting, then increased competition may actually make for a more concentrated set of *a priori* important topics on the agenda. It is not clear however, that traditional news outlets can maintain forever their privileged status in this regard. Some new entrants (e.g. the Huffington Post) have managed to build a relatively strong 'voice' over the last few years.

We also explore what happens when the success of news is endogenous, i.e. if the act of publishing a topic ends-up increasing its likely success. Interestingly, we find that an excessively strong contest tends to concentrate reporting on topics with the highest *a priori* success probabilities. We also find that the number of competitors has a somewhat ambiguous effect on the outcome. If there are too few or too many competing firms then, again agenda setting tends to remain conservative in the sense of focusing on the *a priori* likely topics. These results also resonate to anecdotal evidence concerning today's industry dynamics.

Our analysis did not consider social welfare. This is hard to do as it is not clear how one measures consumer surplus in the context of news. Indeed, the model is silent as to what is consumers' (i.e. readers') utility when it comes to the diversity of news. While policy makers generally consider the diversity of news as a desirable outcome, a view that often guides policy and regulatory choices, it is not entirely clear that more diversity is always good for consumers. As mentioned in the introduction, the media does have an agenda setting role and it is hard to argue that every topic equally represented in the news is a useful agenda to coordinate collective social decisions (e.g. elections). Nevertheless, our goal was to identify the competitive forces that may play a role in determining the diversity of news. Our analysis indicates that these forces do not necessarily have a straightforward impact on diversity.

Our framework can be extended in a number of directions. So far, we assumed a static model, one where repeated contests are entirely independent. One could also study the industry with repeated contests between media firms where an assumption is made on how success in a period may influence the reward or the predictive power of a medium in the next period. A similar, setup is studied with a Markovian model by Ofek and Sarvary (2003) to describe industry evolution for the hi-tech product categories. Finally, our paper generated a number of hypotheses that would be interesting to verify in future empirical research.

References

Cohen, Bernard. 1963. *The Press and Foreign Policy*. Princeton University Press, NJ.

Deighton, John, Leora Kornfeld. 2010. Case study: United breaks guitars. *Harvard Business School* .

Dellarocas, Chrysanthos, Zsolt Katona, William Rand. 2013. Media, aggregators and the link economy: Strategic hyperlink formation in content networks. *Management Science* forthcoming.

Eisensee, Thomas, David Stromberg. 2007. News floods, news droughts, and u.s. disaster relief. *Quarterly Journal of Economics* **122**(2).

Franceschelli, Ignacio. 2011. When the ink is gone: The impact of the Internet on news coverage. *Working Paper* .

Gentzkow, Matthew, Jesse Shapiro. 2010. What drives media slant? evidence from u.s. daily newspapers. *Econometrica* **78** 35–71.

George, Lisa M., Felix Oberholzer-Gee. 2011. Diversity in local television news. *Working Paper* .

Larcinese, Puglisi Riccardo, Valentino, James M. Snyder. 2011. Partisan bias in economic news: Evidence on the agenda-setting behavior of u.s. newspapers. *Journal of Public Economics* **95** 1178–89.

Mayzlin, Dina, Hema Yoganarasimhan. 2012. Link to success: How blogs build an audience buy monitoring rivals. *Management Science* **58**(9) 1651–68.

McCombs, Maxwell. 2004. *Setting the Agenda: The Mass Media and Public Opinion*. Blackwell, Malden, MA.

McCombs, Maxwell, Jian-Hua Zhu. 1995. Capacity, diversity, and volatility of the public agenda: Trends from 1954 to 1994. *The Public Opinion Quarterly* **59**(4) 495–525.

Mullainathan, Sendhil, Andrei Shleifer. 2005. The market for news. *American Economic Review* **95** 1031–53.

Ofek, Elie, Miklos Sarvary. 2003. R&D, marketing and the success of next-generation products. *Marketing Science* **22**(3) 355–70.

Ottaviani, Marco, Peter Norman Sorensen. 2006. The Strategy of Professional Forecasters. *Journal of Financial Economics* **81**(10) 441–466.

Prat, Andrea, David Stromberg. 2013. The political economy of mass media. *Journal of Political Economy* forthcoming.

Snyder, James M., David Stromberg. 2010. Press coverage and political accountability. *Journal of Political Economy* **118**(2).

Stromberg, David. 2004. Mass media competition, political competition and public policy. *Review of Economic Studies* **71**(1) 265–84.

Xiang, Yi, Miklos Sarvary. 2007. News consumption and media bias. *Marketing Science* **26**(5) 611–28.

Xiang, Yi, David A. Soberman. 2013. Consumer favorites and the design of news. *Management Science* forthcoming.

Yildirim, Tuba Pinar, Esther Gal-Or, Tansev Geylani. 2013. User-generated content in news media. *Management Science* forthcoming.

8 Appendix

Proof of Proposition 1: Throughout the proof, we drop the superscript (N) and q_k denotes the probability that any given site out of N players chooses topic k in the symmetric mixed strategy equilibrium. Let k_1, k_2, k_{K+} denote the set of topics that players pick with a positive probability in a potential (symmetric) equilibrium in such an order that $v_{k_1} \geq v_{k_2} \geq \ldots \geq v_{k_{K+}}$. We calculate the expected payoff of firm i when choosing topic k_1 as

$$\mathbf{E}\,\pi_i^{(k_1)} = \sum_{l=1}^{K} \gamma_\ell \,\mathrm{Pr}(\text{topic } k_1 \text{ is a success among } \ell \text{ topics}) \sum_{m=1}^{N} \frac{1}{m} \,\mathrm{Pr}(m{-}1 \text{ other players choose topic } k_1) \tag{5}$$

The first probability is simply $\sum_{k_1 \in S, |S|=l} P_S$, therefore

$$\sum_{l=1}^{K} \gamma_\ell \,\mathrm{Pr}(\text{topic } k_1 \text{ is a success among } \ell \text{ topics}) = p_{k_1} - \sum_{\{k_1 \in S,\ |S| \geq 2\}} (1 - \gamma_{|S|}) P_S = v_{k_1}$$

The second second probability in (5) can be written as

$$\sum_{n_{k_2}+\ldots+n_{k_{K+}}=N-m} \frac{(N-1)!}{(m-1)! n_{k_2}! n_{k_3}! \ldots n_{k_{K+}}!} q_{k_1}^{m-1} q_{k_2}^{n_{k_2}} q_{k_3}^{n_{k_3}} \ldots q_{k_{K+}}^{n_{k_{K+}}}.$$

Since $q_{k_2} + \ldots + q_{k_{K+}} = 1 - q_{k_1}$, we can write the above as

$$\frac{(N-1)!}{(m-1)!(N-m)!} q_{k_1}^{m-1} \sum_{n_{k_2}+\ldots+n_{k_{K+}}=N-m} \frac{(N-m)!}{n_{k_2}! n_{k_3}! \ldots n_{k_{K+}}!} q_{k_2}^{n_{k_2}} q_{k_3}^{n_{k_3}} \ldots q_{k_{K+}}^{n_{k_{K+}}} =$$

$$= \frac{(N-1)!}{(m-1)!(N-m)!} q_{k_1}^{m-1} (1 - q_{k_1})^{N-m}. \tag{6}$$

Therefore,

$$\mathbf{E}\,\pi_i^{(k_1)} = v_{k_1} \sum_{m=1}^{N} \beta_m \frac{(N-1)!}{(m-1)!(N-m)!} q_{k_1}^{m-1} (1 - q_{k_1})^{N-m}. \tag{7}$$

Notice that the above sum is an expectation of β_m, where $m-1$ is distributed Binomially. Let $X(n,q) \sim \mathrm{Binom}(n,q)$ be a random variable and let $G_N(q) = \mathbf{E}\,\beta_{1+X(N-1,q)}$. It is clear that $\mathbf{E}\,\pi_i^{(k)} = v_k G_N(q)$ for any other k topic. Following from its definition $G_N(q)$ is a decreasing, continuous function on $[0,1]$ with $G_N(0) = 1$ and $G_N(1) = \beta_N$ for any $N \geq 1$. Furhtermore $G_N(q)$ is decreasing in N for any fixed q value. An equilibrium has to satisfy

$$v_{k_1} G_N(q_{k_1}) = v_{k_2} G_N(q_{k_2}) = \ldots = v_{k_{K^+}} G_N(q_{k_{K^+}}) \tag{8}$$

and $v_{k_1} G_N(q_{k_1}) > v_k G_N(q_k)$ for any other k topic. Since $G_N(q)$ is decreasing, it follows if players put a positive probability on a topic then they have to put a positive probability on all other topics with higher values. Otherwise deviating to the higher value topic would be profitable. That is, the set $\{k_1, k_2, \ldots, k_{K^+}\}$ has to be of the form $\{O(1), O(2), \ldots, O(K^+)\}$, a set of the K^+ highest value topics. Together with (8), this implies *Part 1* as $G_N(q)$ is decreasing.

For *Part 2*, let $\overline{K_N}$ denote the highest integer for which (8) has a solution with positive q_k iff $v_{k_{\overline{K_N}}} N > v_{k_1}$. The solution with a given set of $\{O(1), O(2), \ldots, O(\overline{K_N})\}$ has to be unique due to the decreasing $G_N(q)$ function. Furthermore, there is no solution, where less than K_N topics receive a positive probability.

To see that K_N is increasing in N, assume on the contrary that $K_{N_1} \leq K_{N_2}$ for some $N_1 > N_2$. Since $G_N(q)$ is decreasing in both q and N, it follows that $q_O^{(N_1)}(j) \leq q_O^{(N_2)}(j)$ for any $j \leq K_{N_2}$. However, then $1 = \sum_{j=1}^{K_{N_1}} q_O^{(N_1)}(j) < \sum_{j=1}^{K_{N_2}} q_O^{(N_2)}(j) = 1$, which is a contradiction.

Finally, for *Part 3*, note that for any k topic that is chosen in equilibrium with positive probability $v_{O(1)} G_N(q_{O(1)}) = v_k G_N(q_k)$. Since $G_N(q)$ falls between β_N and 1, this implies $v_k > v_{O(1)} \beta_N$. $\square$

Proof of Proposition 2:

Let us define $f_\beta(x) = \lim_{n\to\infty} \frac{\beta_{nx}}{\beta_n}$ for $x \geq 1$. The definition implies that $f_\beta(x_1 x_2) = f_\beta(x_1) f_\beta(x_2)$ for any $x_1 \geq 1, x_2 \geq 1$. Also it is clear that $f_\beta(1) = 1$ and that $f_\beta(.)$ is decreasing. Let $x' = 2^{(a/b)}$, where a, b are integers so that $\log_2(x')$ is rational. Then $f_\beta(x') = f_\beta(2)^{(a/b)}$. If $f_\beta(2) = 0$, this implies $f_\beta(x') = 0$. Since $f_\beta()$ is decreasing, $f_\beta(x) = 0$ for any $x > 0$. If $f_\beta(e) > 0$, then $f_\beta(x') = (2^{-\log_2(f_\beta(2))})^{(a/b)} = (2^{(a/b)})^{-\log_2(f_\beta(2))} = x^{-r}$. Since $f_\beta()$ is decreasing and is x^{-r} on a dense set, $f(x) = x^{-r}$ has to hold for any $x > 1$. Therefore, for a monotone decreasing β_n series the $f_\beta(x)$ function can be either $f_\beta(x) = 1/x^r$ with a non-negative r, or a discontinuous function with $f_\beta(1) = 1$ and $f_\beta(x) = 0$ for any $x > 1$, corresponding to $r(\beta) = \infty$.

Recall that in equilibrium (8) holds for every topic with high enough value. As $N \to \infty$, we get $G_N(q)/(\beta_{Nq}) \to 1$ due to the central limit theorem. Also, when $\beta_N \to 0$, we also have $G_N(1) \to 0$, that is, for a large enough N, each topic is chosen with positive probability. Therefore, for any two topics $\beta_{Nq_{k_1}^{(N)}}/\beta_{Nq_{k_2}^{(N)}} \to G_N(q_{k_1}^{(N)})/G_N(q_{k_1}^{(N)}) = v_{k_2}/v_{k_1}$. Since the $f_\beta(x)$ limit exists for any x and is decreasing in x, $q_{k_1}^{(N)}/q_{k_2}^{(N)}$ converges for any pair of topics. The sum of all q values is 1, hence each $q_k^{(N)}$ probability converges to a limit denoted by q_k^∞. When $r(\beta) > 0$, the definition of the $f_\beta()$ function implies $v_{k_2}/v_{k_1} = f_\beta(q_{k_1}^\infty/q_{k_2}^\infty) = (q_{k_2}^\infty)^{r(\beta)}/(q_{k_1}^\infty)^{r(\beta)}$. Since the sum of probabilities is 1, we get the stated results for any $r(\beta) > 0$, including $r(\beta) = \infty$. When $r(\beta) = 0$, all probabilities converge to 0 except one and that has to belong to the highest value topic and converge to 1. $\square$

Proof of Corollary 1:

Let us assume w.l.o.g that $\frac{v_K}{s_K} \geq \frac{v_{K+1}}{s_{K+1}}$. Then, we have $\frac{v_K+v_{K+1}}{s_K+s_{K+1}} \geq \frac{v_{K+1}}{s_{K+1}}$. Note that when $\beta_n = 1/n$, we have $G_N(q) = \frac{1-(1-q)^N}{qN}$ in (8), hence for any k topic,

$$\frac{v_k}{s_k}\left(1-(1-s_k)^N\right) = \frac{v_{K+1}}{r_{K+1}}\left(1-(1-s_{K+1})^N\right) < \frac{v_K+v_{K+1}}{s_K+s_{K+1}}\left(1-(1-s_K-s_{K+1})^N\right).$$

If we had $q_K \leq s_K + s_{K+1}$, then this would yield

$$\frac{v_k}{s_k}\left(1-(1-s_k)^N\right) < \frac{v_K}{q_K}\left(1-(1-q_K)^N\right) = \frac{v_k}{q_k}\left(1-(1-q_k)^N\right)$$

which would lead to $q_k < s_k$, which is a contradiction, since the sum of the q probabilities is 1, just as the sum of the s probabilities $\square$

Proof of Proposition 3: Let us first examine the decision of branded providers, assuming that non-branded providers mix between topics $k_1, k_2, ..., k_{K^+}$. If their proportion is small enough, a branded provider cannot be indifferent between any two topics in the above set, as their profit function is that same as non-branded providers' except for a λp_i term. This additional term leads them to choose topic 1, the topic with the highest prior probability. Given the strategies of the branded providers, we can derive the mixing strategies of the non-branded providers. Without loss of generality, assume that $k_1 = 1$. Following the same lines as in the proof of Proposition 1, similarly to (8) we get that

$$v_1 \tilde{G}_N(q_1^{NB}) = v_{k_2} G_{(1-\alpha)N}(q_{k_2}^{NB}) = \ldots = v_{k_{K+}} G_{(1-\alpha)N}(q_{k_{K+}}^{NB}) \tag{9}$$

and $v_{k_1} G_N(q_{k_1}) > v_k G_N(q_k)$ for any other k topic. $\tilde{G}_N()$ above is defined similarly to $G_N()$, as $\tilde{G}_N(q) = \mathbf{E}\,\beta_{\alpha N+1+X((1-\alpha)N-1,q)}$. Therefore, $\tilde{G}_N(q) < G_{(1-\alpha)N}(q)$, which implies that the solutions q_k^{NB} have to satisfy $q_1^{NB} < q_1$ and $q_k^{NB} > q_k$ for any other k, where q_k's are the solution of

$$v_1 G_{(1-\alpha)N}(q_1) = v_{k_2} G_{(1-\alpha)N}(q_{k_2}) = \ldots = v_{k_{K+}} G_(1-\alpha)N(q_{k_{K+}}) \tag{10}$$

□

Proof of Corollary 2: Following the same steps as in the proof of Proposition 2, we obtain for any two topics other than topic 1 that $\beta_{(1-\alpha)Nq_{k_1}^{(N)}}/\beta_{(1-\alpha)Nq_{k_2}^{(N)}} \to G_{(1-\alpha)N}(q_{k_1}^{(N)})/G_{(1-\alpha)N}(q_{k_1}^{(N)}) = v_{k_2}/v_{k_1}$. For topic 1 and an arbitrary topic k, we obtain $\beta_{\alpha N+(1-\alpha)Nq_1^{(N)}}/\beta_{(1-\alpha)Nq_k^{(N)}} \to v_k/v_1$. These yield $v_{k_2}/v_{k_1} = (q_{k_2}^{\infty})^{r(\beta)}/(q_{k_1}^{\infty})^{r(\beta)}$ and $v_k/v_1 = (q_k^{\infty})^{r(\beta)}/(q_1^{\infty} + \alpha/(1-\alpha))^{r(\beta)}$, .

□

Proof of Proposition 4: High-type players clearly always chose the topic that becomes successful as they are able to predict perfectly which one it will be. The equilibrium strategies for low-type players can be determined along the the same lines as in the proof of Proposition 1. Equation (8) in this case is transformed to

$$v_{k_1} \hat{G}_{(N}(q_{k_1}^L) = v_{k_2} \hat{G}_N(q_{k_2}^L) = \ldots = v_{k_{K+}} \hat{G}_{(N}(q_{k_{K+}}^L) \tag{11}$$

where $\hat{G}_N(q) = \mathbf{E}\,\beta_{\mu N+1+X((1-\mu)N-1,q)}$. Since $\hat{G}_N(0) = \beta_{\mu N+1}$ and $\hat{G}_N(0) = \beta_N$, any k topic chosen with a positive probability has to satisfy $v_k > v_{O(1)}\beta_N/\beta_{\mu N+1}$. □

Proof of Corollary 3: We can show that all mixing probabilities converge as in the proof of Proposition 2. Furthermore, we obtain for any two topics that $\beta_{\mu N+(1-\mu)Nq_{k_1}^{(N)}}/\beta_{\mu N+(1-\mu)Nq_{k_2}^{(N)}} \rightarrow v_{k_2}/v_{k_1}$. This yields $v_{k_2}/v_{k_1} = \left(\frac{\mu+(1-\mu)q_{k_2}^{\infty}}{\mu+(1-\mu)q_{k_2}^{\infty}}\right)^{r(\beta)}$ if both $q_{k_1}^{\infty}$ and $q_{k_2}^{\infty}$ are positive. Clearly, this equation cannot hold for too small v_k values, giving the threshold stated in the proposition. Topics with values below the threshold cannot receive a mix probability converging to a positive number. For topics with a positive mix probability in infinity, the previous equation implies the stated formula by summing all positive probabilities to 1. □

Proof of Lemma 1: Player i's profit when all players choose topic k is $\pi_i^{(k)} = (1/N^r)(\lambda + (1-\lambda)v_k)$. When player i deviates to topic k', the profit becomes $\pi_i^{(k)} = \lambda/N + (1-\lambda)v_1$. Hence the most profitable deviation is to topic $O(1)$ if $k \neq O(1)$ and to topic $O(2)$ if $k = O(1)$. Therefore $k \neq O(1)$ can be an equilibrium topic choice for all players iff

$$D(k,r,N,\lambda) \stackrel{def}{=} \frac{1}{N^r}(\lambda + (1-\lambda)v_k) - \frac{\lambda}{N} - (1-\lambda)v_1 \geq 0 \tag{12}$$

Clearly, $D(k,r,N,\lambda)$ is decreasing in k, proving the lemma. □

Proof of Proposition 5: Checking that $D(k,r,N,\lambda)$ from (12) increases in λ proves that $\bar{j}(\lambda, N)$ is also increasing in λ.

For *Part 1*, let $k \neq O(1)$ be an arbitrary, but not the highest value topic. Since $D(k,r,1,\lambda) = (1-\lambda)v_k < 0$ and $D(k,r,N,\lambda) \underset{N\to\infty}{\rightarrow} -(1-\lambda)v_1 < 0$, topic k cannot be a pure equilibrium for small or large N's. Furthermore,

$$\frac{\partial D(k,r,N,\lambda)}{\partial N} = \frac{1}{N^2}(\lambda - r(\lambda + (1-\lambda)v_k)N^{1-r}),$$

which is negative for small N's and positive a certain N^* value when $r > 1$ (and all negative when $r = 1$). Therefore, $D(k,r,1,\lambda) < 0$ when $r \geq 1$, proving *Part 1*.

When $r < 1$, the derivative is positive for small N's and negative for large N's with a single root at $N^* = \left(\frac{\lambda}{r(\lambda+(1-\lambda)v_k)}\right)^{1/(1-r)}$. Plugging this value in shows that $D(k,r,N^*,\lambda) > 0$ for high enough values of λ, proving *Part 2*.

Finally, since $D(k, r, N, \lambda)$ is inverse-U shaped in N, being negative at $N = 1$ and $N \to \infty$, with a positive maximum for high λ's, there must be a $1 < \underline{N} \leq \overline{N}$ interval where it is positive, proving *Part 3*. $\square$