Strategic grouping and search for quality journalism, online versus offline

Matthew Ellman[*†]   Tomás Rodríguez[*‡]

October 2016

First Public Draft

Abstract

This paper investigates how supply-side factors influence the search for quality content in online and offline environments. We show that lower fixed costs of online publishing reduce the incentives to bundle content, as compared to offline journalism. In the presence of asymmetric information over journalistic quality, bundling of content by journalists who publish as a group generates positive informational externalities for users. Journalists group assortatively, better journalists having better partners. Then a consumer who discovers one quality journalist, has found several. The online environment, by reducing the pressure to group up, can lower welfare in our baseline model. We establish conditions for this result and investigate a number of countervailing forces.

Keywords: Media economics, quality, search, links, matching.
JEL Classifications: L13, L82.

[*]We gratefully acknowledge financial support from the Net Institute, http://www.Netinst.org.
[†]Institute for Economic Analysis (CSIC) and BGSE. E-mail matthew.ellman@iae.csic.es.
[‡]Universitat Autónoma de Barcelona, MOVE and BGSE. E-mail: tomas.rodriguez@uab.cat.

1 Introduction

Digital optimists predicted that dramatic reductions in the cost of self-publishing would "democratize" journalism, enabling new journalists to publicize their insights directly to a wider public, without relying on incumbent media companies or financial support. Less passionate scholars from the economics profession also predicted greater specialization of outlets, permitting a tighter match between content offerings and people's interests. Both views suggested a reduction in the concentration of media consumption online. However, despite the mass of new reporting and commentary, the news market on the internet remains highly concentrated into a small set of well-established media outlets. We argue that people's difficulties in distinguishing the wheat from the chaff among journalistic content can explain why readers may fall back on these established outlets even when people find their quality quite low, or dislike their ideological positions.[1] The puzzle is to understand why online intermediaries, like search engines, that have dramatically reduced the cost of search in so many markets, and *do* work well for finding specific news topics, do not also serve well for filtering out poor quality content, the chaff, in the case of news and analysis.[2]

We adopt a supply-side analysis to help understand why the open-journalism hopes of digital optimists have so far proved elusive. We investigate peer review as a form of quality assurance, but in a context where there are no centralized bodies offering certification services, since news is notoriously subjective and hard to verify. Instead we look for decentralized mechanisms that permit distributed actors, the journalists, to convey information via their organizational choices.

We then seek to contribute to the understanding of how a substantial reduction in the fixed costs of operating a viable news outlet, associated with digital technologies and the world wide web, is likely to affect the size and quality of news organizations and consumer welfare. The central idea of our model is that, given asymmetric information, when journalists sort into outlets according to their qualities, it becomes easier for consumers to find high quality articles, because once they find one, they find several. Journalists do not take this increase in consumer welfare into account when deciding whether to form a group. As a result, fewer groups tend to form than would be socially optimal. They fail to internalize all the user benefits from grouping together, because users pay

[1] Quality may refer to any characteristic that makes a reader value an article and for which many readers have a similar valuation. But we emphasize informational accuracy as the key component of quality because this is the type of quality that we believe is hardest for readers, and in consequence search engines, to evaluate.

[2] It is well known that selling information is difficult. Information has some credence good properties and, while quality can be learned by inspection and experience to some degree, it is hard for the average reader to judge how well-informed is a journalist writing on, say, Syria. We read experts precisely because we do not know as much as they do. Our contribution is to investigate the channels through which the reduced costs of publishing online affect the difficulty readers face in their quest for high quality content.

nothing directly and cannot easily be charged indirectly via advertising as a nuisance.[3] In the printed news environment, the relatively high fixed costs of operating a news outlet partly mitigate the tendency to form groups too rarely. Removing these high fixed costs diminishes journalists' incentives to form large groups and we show that this can worsen the lot of consumers.

In other settings where assortative grouping reduces search costs, the incentives for one set of actors to group together need not depend on how consumers expect the overall set of actors to organize themselves. A prime example is the retail sector, in which commercial centers are pervasive and often concentrate stores of similar quality. Since consumers can readily inspect whether the additional stores in the commercial center are to their liking, and decide to buy additional products accordingly, consumers' beliefs about the grouping or co-location strategies of retailers in general have far less impact on the profits that a given group of retailers would gain from grouping together.

By contrast, whether a journalist of a given quality is able to direct traffic to other members of her outlet depends on the beliefs held by readers. Since inspecting an article is very costly, the existence of a bundle of high quality articles is not useful to the reader by itself. It becomes useful only to the extent that the aggregate grouping pattern gives rise to beliefs that allow him to make the correct decision over whether to read additional articles in an outlet based on his initial experience. This feature of our setting makes it quite different from other studies of markets with informational asymmetries where grouping creates positive externalities.

We show that the lower fixed costs of operating news outlets in the online world leads to substantial separatism, that is a reduction in the size of groups, and we demonstrate that this tendency has important consequences for the average reliability of the content that is read online as compared to reading outcomes in a world with only offline journalism. In brief, we prove that good journalists will tend to group together and that this positive assortative matching (PAM), or indeed a weaker variant of sorting that we call threshold assortative matching (TAM), reduces readers' costs of finding quality content. To repeat the basic intuition: when they find one good article or journalist, they have found many. Then we show that shifts that reduce the net benefits of grouping in the online setting make grouping less likely and this translates into a greater challenge for readers seeking to find good content. As a result, the average quality of news consumed on the web can fall. More generally, the analysis points out some new factors that may explain why we have, by and large, failed to see fulfillment of the web's promise to give voice to the many undiscovered but high quality individuals who were effectively silenced in the pre-internet world.

[3] The direct price on users is zero as is common in such two-sided markets. Some online outlets do use pay-walls, but by and large, the revenues of online and, now even many offline, news outlets stem from advertising.

Surprisingly, we also find that negative assorting is possible for some parameter ranges. This cannot be excluded by the intuitive criterion, except when diseconomies of scale are relatively high, in which case it does deliver unique positive assorting. Other refinements may rule out these less intuitive equilibria, but we also infer that there is much to be gained by studying a richer model with journalist specialization. This richer model would distinguish between journalists who write original articles and journalists who play a more curatorial role, offering content digests and recommendations. In such a setting, we may find that less skilled writers specialize as conduits for sending traffic to better writers and investigators.

Literature Review

The main contribution of this paper is to study how the organizational choices, over how and if to pair up, of a decentralized set of actors, journalists, convey information to a less informed set of actors, called users. The users "purchase" media content which is partly an inspection and partly an experience good, but they pay with their attention which is a scarce resource, and journalists gain revenues from advertisers who want to reach the users. For focus, we assume that the impact of advertising on users is neutral on net, with the gains from finding valued products just compensated by the nuisance of advertising.

Given the focus on pairing choices, the paper overlaps methodologically with the field of matching, but at the same time, it shares an applied focus with the literature on reputation and screening. We begin by discussing prior work on matching and our theoretical novelty there. We then discuss briefly the literature on reputation and screening, before turning to recent work on aggregators and links in the internet environment.

A large literature on matching studies how actors match into groups, usually pairs and predominantly assuming two sides, such as firm/worker or male/female or hospital/doctor. As we study matching among peers, the "one-sided pairing" variant of matching is relevant. We refer to Legros and Newman (2002) for a thorough analysis of monotonicity in the case with an exogenous profit function, especially as their paper is one of the few that treats a one-sided pairing environment. See also Legros and Newman (2007) for an in-depth study of the problem with a continuum of agents, but in the context of two-sided matching.

In the matching literature, the value of a group, or coalition, is specified exogenously as some function of the characteristics of the members of the group. However, in our setting, the value of a group also depends on user beliefs, because users observe a signal about the group before they decide what actions to take. In work on matching with adverse selection, the uninformed party moves first, committing to a contract before any matching takes place. The uninformed party's beliefs about matching determine the proposed menu of contracts, but at the stage of group formation, the values of each possible group (of informed parties) are pinned down by those contracts; the values do not depend on any inferences that the uninformed party's might make after observing what the informed parties do since the uninformed parties have no further move to make. By contrast, in our setting with many uncoordinated users, as uninformed parties, it is impossible for them to commit to reading strategies before journalists choose how to organize production. So we contribute to the matching literature by studying a context where group values naturally depend on beliefs about matching by a set of uninformed agents (the users) who take actions after the group matching stage.

Few papers consider asymmetric information in a matching context, but there are one or two exceptions that contemplate asymmetric information between the matches. Liu et al. (2014) is an important example. This paper provides a rigorous treatment of asymmetric information, but has no asymmetrically uninformed side, corresponding to our users. Instead, the actors who match up into pairs do not know each others' qualities. This introduces a range of important complications for which they develop useful techniques, but which cannot serve in our context since the type inferences affecting matching strategies become irrelevant after the matching stage.

We close the discussion of matching by mentioning an example of applied work where matching into pairs is used to reduce informational problems: Ghatak (1999) and Ghatak (2000) study a model of group lending with asymmetric information in which borrowers have information that lenders lack. They show that group lending schemes imposing joint liability can induce peer selection into groups that allows lenders to perform better screening. These papers exemplify the above claim that prior work has assumed that the uninformed party is a single actor who moves first, with full commitment, avoiding the need to model a strategic reaction after the matching stage. For a more recent example of this type of model and closer to our digital context, see **?**'s study of interpersonal bundling.

Moving on to the applied objective of our analysis, there is a large literature on the potentials and risks of online reputation mechanisms aimed at reducing moral hazard. One branch studies how, by facilitating elicitation and publicity of consumer feedback, digital technologies, using the internet as communication channel, can enhance quality assurance. This reputation literature has focused on standard goods markets and not the fundamentally two-sided market for media, but the more important difference is that we study the role for peer evaluation, rather than buyers evaluating sellers. For examples of work on reputation-building via online mediated feedback, see Bar-Isaac and Tadelis (2008); Cabral and Hortacsu (2010); Dellarocas et al. (2013). In this literature, the role of sellers evaluating each other has received little treatment. Certification has been studied, but we work in the context of distributed and decentralized sellers. Two closer papers are those by Dellarocas et al. (2013) and Jeon et al. (2012), but still in both these works, the party that links to the others is a centralized party, called the aggregator, rather than decentralized peers. Those papers are therefore closer to studies of recommender sites, where the intermediary that recommends what a user should look at. Burguet et al. (2015) bears mention since it studies how distortions in search results for content, interact with the distortions in search results for products. Nonetheless, Jeon et al. (2012)'s model of audience-expansion and business-stealing are close precursors to some of the effects in our paper. An important point is that actors, such as aggregators, might link to the best journalism available. We build on Burguet et al. (2015) to show why this may not obtain for new media content, even when link-contingent transfers are feasible. There is also a nascent empirical literature on hyperlinking in online news that is highly relevant to our application; see for instance George and Hogendorn (2012). Their empirical results point to an important role for aggregators in enabling people to find smaller media outlets. Our message is that aggregators and other online intermediaries are valuable for finding topics of interest but are not sufficiently effective at filtering for quality to prevent the concentration of media consumption on a handfull of major and well-established outlets.

2 Model

A unit mass of users, each with a demand for two news articles, value quality content but do not know which journalists have the highest quality.

The default option. Users can always access a set of established media outlets. We call the best option from this set of established journalism, their "default." Users know that this default offers them at least two (relevant) articles of quality z_0.

Motivating remarks. These default outlets may serve many people who do not wish to risk trying to find better quality, but we do not model those people. Instead, we study users who would like to discover new journalists. These users might be relatively dissatisfied with established media options on grounds of quality, style, bias, or topic coverage, or they might just enjoy being exposed to new voices and original or non-mainstream viewpoints. Trying out new journalists involves a risk, because articles are largely experience goods (with a dose of credence good properties) and evidence of their quality only slowly percolates into the public realm. Quality signals can be very noisy, especially for subjective topics and political issues where different users have different personal biases; as explained below, search engines, aggregators and online social networks learn from the subjective views of their users and journalists' linking strategies.

We model the challenge for users to find quality content in the face of imperfect information and we focus on how their search for quality interacts with the strategic grouping and linking behavior of the unestablished journalists. We abstract from the strategic choices of default media; the default outlets are not active players and when we refer to journalists, we mean unestablished journalists. In brief, we want to understand how the interactions between a decentralized set of journalist producers can facilitate users' search for quality. Journalists are necessarily decentralized because they can at most form dyads. Some upper bound on organizational size is natural. Antitrust prevents huge coalitions; new journalists naturally start off small; even though conceivable, a state monopoly that might coordinate the journalists would typically not be trusted in countries that wish to fiercely protect media independence from state interference.

Journalists. A set J of journalists produce news content. Each $j \in J$ produces one news article of quality z_j. These qualities are *i.i.d.* draws from a differentiable distribution $F(\cdot)$ on $[0,1]$. We sometimes use the uniform distribution $U[0,1]$ to provide simple analytical solutions. We let $\mathbf{z}$ denote the size J vector of realized qualities.

Matching. In stage 0, journalists either form organizations or stay out of the market. Each j can pay a cost c_m to publish as a monad, or with her pair, she must cover the cost c_d to pair with a peer j' as a dyad. If she pays nothing, she simply stay out of the news market and her payoff is zero.

To reduce the number of subscripts, we let $c = c_m$ and we define the per-capita added cost of dyads, $\Delta = (c_d - 2c_m)/2$. If costs are paid for equally by organizational members, j pays c in a monad and pays $c + \Delta$ in a dyad.

The mapping $m : J \to J \cup \{\varnothing\}$ denotes the choices for each journalist, j. That is, $m(j)$ describes j's pairing outcome in stage 0: $m(j) = j$ if j forms a monad, (j), $m(j) = j' \neq j$ if j pairs with $j' \in J \setminus \{j\}$ to form a dyad, (j, j'), and $m(j) = \varnothing$ if j stays out of the market. All journalists observe each others' qualities (though users do not) and pairing requires the agreement of both journalists in the proposed pair. Notice that $m(m(j)) = j, \forall j$, because pairings are undirected links.

Remark on pairings and links. In the extended model, we will distinguish between organizational pairings and "referral" links (such as hyperlinks in the online case), by calling only the latter, links. In the baseline model, we do not distinguish between pairings and (undirected) links, because we do not allow for any other links and we assume that any two journalists in a dyad are automatically linked to one another. This turns out to be the optimal strategy for dyads, so nothing is lost in making this simplification, as we prove later. Moreover, links, as referrals, are arguably automatic in the offline setting where two journalists form a dyad to publish their articles in a common newspaper. By bundling their content in the same physical vehicle, their articles automatically become linked in the sense that a user who reaches one article can readily switch to the other journalist's article. Online, we suppose that the journalists in a dyad share a website or group affiliation and make this structure transparent to users, so that they can readily learn of the existence of the partner of a journalist in a dyad.[4]

Users. Each of the unit mass of users starts his search for two articles by drawing a signal. To capture the increasing cost of search over time, or scarcity of recommendations from trusted friends, we assume that each user can only draw one signal (see Section 4.1 for a discussion). This signal identifies an article by a journalist j, which we call that user's "primary" article and "primary" journalist; in the baseline, it is equivalent to speak of journalists and articles, since they are bijectively related. Whether or not he reads j's article, the user can always switch to a default article and if his primary journalist j links to another journalist j', as arises for all dyads, he can also switch to that journalist's article, which we call his "secondary" article (and j' is his secondary journalist). Users' signals are distributed across articles according to the journalists' "prominence," described just below. In the baseline model, each signal identifies an article by some journalist j and fully reveals its quality z_j to the user.[5] Each user knows the distribution of all articles, but only observes the quality of his primary article.

In sum, he learns j's quality z_j and he also learns how to costlessly reach j's article, and the article of j's pair $m(j)$ whenever j is in a dyad, since he observes the pairing status and out-links of his primary journalist j.

Article prominence. Users are sent to articles according to their prominence, $p(j, m(j); \{m(k)\}_{k \in J}, \mathbf{z})$, so more "prominent" journalists win more primary users. We call these users their primary or direct traffic. We think of intermediaries such as search engines, social contacts, aggregators or other recommenders as determining this prominence function. This is exogenous in the baseline model but we motivate our specific functional form as follows.

First, we let prominence depend on a journalist's pairing status ($d_j = 1$ for dyads, $d_j = 0$ if j is in a monad), which we take to be readily observed, but not on the journalist's quality z_j. This reflects the extreme case where nobody has yet inspected any of these journalists' qualities yet. Implicitly, intermediaries may observe signals of how closely a given article fits the given user's interests, and recommend it or not accordingly. Even though we are supposing that intermediaries do not observe quality directly, they might take account of the facts that journalists who self-select to form dyads may a) have higher quality (Google's PageRank promotes nodes that have greater network centrality and dyads are always more central than monads in our baseline) and b) benefit users by

[4]Abstracting from editorial design, specific cross-referencing, the front page, contents listing and subtle layout effects, these links are necessarily undirected. So this is a reasonable starting point, even though it is vital to also allow for asymmetries. We do this when we consider directed links; they are especially important in the online setting.

[5]It simplifies to assume either a perfectly accurate signal of the primary article, or the opposite extreme assumption, that the user only learns its quality by reading it, as with experience goods. The experience good view generates somewhat different results to the baseline inspection model, but we will treat both and prove that our key results apply in either case.

offering a link to the article's pair.

Second, we assume symmetry across articles in terms of user relevance probabilities. So each journalist j receives a visit from a fraction p of the unit mass of users if j is in a monad ($m(j) = j$) and a fraction σp if j is in a dyad ($m(j) \neq j$), where σ is a function of the number of monads M and dyads D in the market.

With Assumption A1 below ensuring that no journalists stay out of the market, M is simply $J - 2D$ and we can express σ as a function of the number of dyads alone. Where relevant, we emphasize this dependence by writing $\sigma(D)$ instead of just σ. In the analysis, we focus on the case where $\sigma \geq 1$ and $\sigma'(D) \leq 0$: dyads stand out but each dyad stands out less prominently if it is among many other dyads. Since we assume each user finds exactly one signal, sending him to exactly one article, the sum of fractions of the population visiting each of the available articles must equal 1. So p is given by:

$$p = \frac{1}{M + 2\sigma(D)D}$$

In sum, each journalist in a monad has primary traffic p and each journalist in a dyad has primary traffic $\sigma(D)p$. Secondary traffic depends on how users react to their signal and primary article; we now describe the user's choices which determine this referred, indirect or "spillover" traffic.

User choice. When the user's primary article is in a monad, he has two options: he either (1) directly reads two articles from the default outlets or (2) reads that article and then switches to his default for his second article.

If instead his search leads to an article in a dyad or pair, he has two alternatives in addition to (1) and (2): (3) he reads both articles from the pair and (4) he reads only the pair of his primary article from the dyad and one article from a default outlet. Notice that he chooses between these four options *after* seeing his first and only signal about the first article's quality.

In the choice of what to read – what "page or site" to visit and whether to read the article there – the user's signal and his primary article's links determine what he knows. In particular, a user arriving on his primary article, learns a perfect signal of its quality z_j, while he reaches his secondary article armed only with the knowledge that his primary article is paired with it.

Timing

Stage 0. Journalists form organizations or stay out of the market.
Stage 1. Each user draws a single signal, identifying one journalist j and her quality z_j. He then uses this information to choose between the courses of action (1), (2), (3) or (4), described above.

User attention and the profit generated by outlets

Whenever a user chooses to read an article, we assume that he devotes an amount of attention that is linear in the article's quality.[6] In addition, we assume that the first two articles the user reads generate a utility for him that is additively separable and proportional to the article qualities. Moreover, we assume that the profit from selling ads to a given user who visits j's site includes a fixed term $\alpha > 0$ and a variable term that is proportional to user attention and hence to the article quality z of any article that he reads.[7] The fixed term $\alpha > 0$ is the ad revenue that results from the possibility to distract any user who arrives at the site.

A visitor to j who knows j's quality perfectly before reading can switch away at no cost to the default media option, giving utility z_0. So he chooses to stay when and only when $h_j = \mathbb{1}\{z_j \geq z_0\}$ equals one: h_j indicates whether j's quality z_j is high enough to merit some attention instead of switching away to the default; we also let $h(z) = \mathbb{1}\{z \geq z_0\}$, giving $h_j = h(z_j)$.[8] We call j, "high" quality or high type if $z_j \geq z_0$ and say that h_j is the indicator for high quality.

A primary visitor to j learns j's quality perfectly before choosing his attention, so his attention and utility gain are given by $h_j z_j$. Suppressing the subscript, primary visits generate utility hz and ad revenue $\alpha + hz$. If j is in a dyad, every primary visitor to j's pair, $j' = m(j) \neq j$ is a potential secondary visitor to j. These (potential) secondary users have no direct signal; recall that recommendations and inspection effort are scarce. Instead, they use $j' = m(j)$'s quality to decide whether to switch to the default or to follow j''s link to j, creating a within-dyad spillover. They follow a link if they have a weak interest in reading the linked content, that is j's article: the condition is $E[z_{m(j')}|z_{j'}] \geq z_0$. This expectation depends on user beliefs about journalists' pairing strategies. When the expectation is high enough, users are sufficiently optimistic about j''s neighbor, they follow the link and generate ad revenue $\alpha + z_j$. Note that there is no high type indicator: only z not hz. On the other hand, defining the "good neighbor" function $g_{j'}$ to indicate whether j' is a good neighbor visitors will want to visit $j = m(j')$, the ad revenue generated from secondary visits is $g_{j'}(\alpha + z_j)$; we say that j' is a good neighbor if $g_{j'} = 1$.

It follows that the revenue to any journalist j and his pair $m(j)$ per primary visitor

[6] This requires some quality awareness and this is consistent with full consciousness, as we now explain. For the user's primary article, we already assumed that the user receives a perfect signal of quality. For the secondary article, the user does not necessarily learn its quality before deciding whether to read. Nonetheless, we suppose he learns it quickly enough to devote quality-proportional attention, yet slow enough that, by the time he realizes quality is low, it is too late for him to switch away.

[7] It is easy to microfound this attention and advertising revenue formulation; the more general idea is simply that rational users engage more when content is beneficial to them, and higher engagement implies more opportunities for journalists to sell ads.

[8] We assume throughout that users read from a new journalist outlet when indifferent between it and the default.

to j's outlet equals:

$$\alpha + h_j z_j + d_j g_j(\alpha + z_{m(j)})$$

Recall that h_j, d_j, g_j are the indicator functions for j being high quality, part of a dyad, and a good neighbor. Given the costs of forming monads and dyads, $c_m = c, c_d = 2(c+\Delta)$, the profit functions for monads and dyads are therefore:

$$\pi(j) = (\alpha + h_j z_j)p_m - c_m \tag{1}$$

$$\pi(j, j') = \big((2\alpha + h_j z_j + h_{j'} z_{j'}) + (g_{j'}(\alpha + z_j) + g_j(\alpha + z_{j'})\big)p_d - c_d \tag{2}$$

As noted earlier, we have a more complicated problem than in standard matching models, because these remarks imply an endogenous payoff of monads and especially dyads, because prominence p depends on how many journalists, D, form dyads and because good neighborliness, which affects dyadic payoffs, depends on users' beliefs. We consider large J to limit journalists' mutual externalities when deviating in pairs or individually.

To make the analysis much simpler to solve, in what follows, we focus on an economy of $k - tuples$, in which there are k copies of each type where k is even and strictly exceeds the group size, 2, so $k \geq 4$. Having $k \geq 3$ ensures that there is competition between journalists for every position in every group. Having k even guarantees the existence of stable equilibria without need for randomization, since it ensures that journalists of any given type can be matched in the same way, including together, without any one of them ever having to be left out.

Stable configurations

The main objective of our analysis is to predict the news organizations that form as a function of the parameters of the model. To do so we assume that in stage 1 users hold beliefs that are consistent with the behavior of journalists in stage 0, and given these beliefs, behave optimally. To begin with, we impose no restriction on the beliefs associated with behaviors that lie off the equilibrium-path.[9] In fact, the optimal behavior of users in stage 1, conditional on users' beliefs, is already reflected in the payoff functions introduced above, it just remains to make sure that these beliefs are consistent with journalists' behaviors. We then use the concept of pairwise stability to predict journalists' pairing choices.[10] A *stable configuration* is a mapping, $\{m(j; \mathbf{z})\}_{j \in J}$, that specifies, for each possible vector $\mathbf{z}$ of types, a pairing on the set of journalists that is feasible and

[9] For example, the beliefs induced by observing a member of a dyad having a quality z which, given the stage 0 behavior of journalists, should only ever be observed in monads.

[10] Pairwise stability is a standard concept used for studying network formation, in settings in which it is preferable not to impose detailed structure on the dynamics of the formation process. As discussed in detail in what follows, in our context pairwise stability

stable. To characterize stable pairings (individually rational and immune to pairwise deviations), we define a bargaining outcome as a vector of payoffs $\mathbf{w}$ where journalist j gets w_j satisfying,

Stability: For all j, j', $w_j + w_{j'} \geq \pi(j, j')$, $w_j \geq \pi(j)$ and $w_{j'} \geq \pi(j')$.

Feasibility: For all $j, w_j + w_{j'} \leq \pi(j, j')$ if $j' = m(j) \neq j$ and $w_j \leq \pi(j)$ if $j = m(j)$.

Notice that this definition of feasibility precludes transfers between journalists who are not linked together.[11]

3 Analysis

We begin by introducing two concepts that allow us to characterize stable configurations and user welfare.

Definition 1 ($\mathbf{TAM}(\hat{z})$). *A pairing satisfies $\hat{z}$-threshold assortative matching if for any vector of types, and all pairs $j, j' = m(j)$, $z_j \geq \hat{z} \iff z_{j'} \geq \hat{z}$.*

If a configuration satisfies $\mathbf{TAM}(\hat{z})$, a user who lands on an article of quality weakly above $\hat{z}$ knows that any linked article is at least of quality $\hat{z}$ as well. Similarly, if he finds an article of quality below $\hat{z}$ then he knows he cannot expect any linked article to have quality $\hat{z}$ or higher. Note that **TAM** is a weak concept; a configuration may enable the user learn much more from the first article's quality. For instance, perfect positive assortative matching within all dyads, which requires $\mathbf{TAM}(\hat{z})$ to hold for all possible cutoffs $\hat{z}$, enables the user to infer that the pair has an identical quality to that of his primary article.

We also define a specific refinement, $\mathbf{TP}(\hat{z})$, of $\mathbf{TAM}(\hat{z})$ that requires all journalists with quality weakly above threshold $\hat{z}$ to form dyads and all journalists with lower quality to form monads or not enter.

Definition 2 ($\mathbf{TP}(\hat{z})$). *A pairing satisfies $\hat{z}$-threshold pairing if for any vector of types, all j with $z_j \geq \hat{z}$ form dyads and all lower quality journalists form monads or stay out.*

[11] We permit arbitrary pairwise deviations since we allow any pair of journalists to form a dyadic production unit. Any theory that contemplates the formation of dyads must implicitly assume that pairs can sometimes solve the problem of coordinating and agreeing to collaborate, so deviations in pairs should also be considered. For this reason, imposing standard pairwise stability is the standard way of modeling the economic processes that determine which pairs will form and create an undirected network. Jackson and Wolinsky (1996) justify this approach for network-theoretic models. Becker (1973) earlier exploited it to study two-sided matching problems. Legros and Newman (2002) provide a general analysis for games with a continuum of players and discuss the one-sided case where actors, like our journalists, match with each other, rather than a distinct set of actors. In game-theoretic terms, we study the core of the cooperative game where coalitions are restricted to have at most two players; our non-cooperative model of the interaction between users and journalists endogenizes the value functions as the payoffs for each possible group (monads and dyads), given users' beliefs about the prior group-formation choices.

We begin by stating and proving a Lemma which we use extensively throughout what follows and which shows that in our economy of $k - tuples$, all agents of the same type must earn the same payoff.

Lemma 1 (Equal treatment). $z_j = z_k \Rightarrow w_j = w_k$.

Proof of Lemma 1:
Without loss of generality suppose that $w_j < w_k$. First, assume that j and k are not part of the same pair. If k is a monad, then $w_j < w_k = \pi(k) = \pi(j)$. If k is in a dyad then $w_j + w_{m(k)} < w_k + w_{m(k)} = \pi(k, m(k)) = \pi(j, m(k))$. Both cases contradict stability. The last equality follows from the fact that $z_j = z_k$ implies $h_j = h_k$ and $g_j = g_k$. Second, suppose that j and k are part of the same dyad, then given at least 3 journalists of the same type (we assumed at least 4), there is some journalist m with $z_m = z_j = z_k$ who is not part of this dyad. Applying the above argument separately to m, j and m, k shows that $w_m = w_j$ and $w_m = w_k$, contradicting $w_j < w_k$. ∎.

We can learn a lot about configuration stability by considering individual rationality and pairwise deviations by twins. Immunity to these deviations, together with equal treatment, nails down the payoff of dyads formed by journalists in the same (h, g) class.

First, consider monads. If j is in a monad, individual rationality requires j to receive $w_j = w_m(j) = [\alpha + h_j z_j]p - c$ and this must weakly exceed the zero payoff from non-entry. The following assumption guarantees that this last condition holds, so that we can neglect non-entry. This simplification is valuable because it lets us focus on quality effects that are driven by matching alone; otherwise, reducing entry costs lowers the average quality of active journalists because the lowest quality journalists gain the least from producing as monads.

Assumption A1: $\alpha \underline{p} \geq c$ *where* $\underline{p} = \min_D (J + 2(\sigma(D) - 1)D)^{-1}$.

The assumption ensures that monads are viable even for the least productive journalist and lowest possible value of $p = \frac{1}{M+2\sigma D}$ taken over all possible realizations of $\mathbf{z}$; $p = 1/(M + 2\sigma D) = 1/\left(J + 2(\sigma(D) - 1)D\right)$. Note that if σ is constant, we have $1/p \in [J, \sigma J]$ but there are good reasons to expect σ to be decreasing in D. One reason is that network centrality rises for connected nodes when other nodes become less connected, and algorithms used in online social networks and search engines use this notion to determine prominence of nodes like webpages or websites, as in the famous case of Google's PageRank. Another reason is that if dyads offer more and often better content to users, recommenders may promote dyads since they are more useful for users. For any given configuration, a pairwise deviation either has no effect on p or it simultaneously shifts both M and $2D$ by 2 units but in opposite directions, changing $1/p$ by at most $2(\sigma - 1)$. We assume J to be large, so that we can neglect changes in p when assessing stability. In the special case where $\sigma = 1$, $\bar{p} = p = 1/J$ and p never changes.

Second, we consider dyads consisting of any two journalists j, j' whose qualities lie on the same side of z_0 (both high or both low), who are equally good or bad neighbors (both $g = 1$ or both $g = 0$) and who are sometimes matched together. Equivalently, consider any journalist j who is sometimes matched with another journalist in the same class, where we define classes by the four possible values of (h, g). The dyadic surplus for any such pair is linear in the sum of their qualities, so it equals half the sum of the two twin-pair surpluses. To see this, note that:

$$\begin{aligned}\pi(j, j') &= \big((2\alpha + h_j z_j + h_{j'} z_{j'}) + (g_{j'}(\alpha + z_j) + g_j(\alpha + z_{j'})\big)p_d - c_d \\ &= \big((2\alpha + h(z_j + z_{j'})) + g(2\alpha + z_j + z_{j'}\big)p_d - c_d\end{aligned}$$

The threat of a twin pair deviation guarantees that each party receives exactly half her twin pairing payoff which we denote,

$$w_d(j) = \frac{\pi(j, \ddot{j})}{2} = \big((\alpha + hz) + g(\alpha + z)\big)\sigma p - c - \Delta \tag{3}$$

If $h = 1$, this simplifies to $w_d(j) = (\alpha + z)(1 + g)\sigma p - c - \Delta$ and if $g = h = 1$, it simplifies to $w_d(j) = (\alpha + z)2\sigma p - c - \Delta$. We state this formally as a Lemma (see Appendix for an explicit proof), adding the monad result from above, but first some terminology: we say that a dyad is "exogamous" or mixed class if its two journalists are from different classes $hg \neq h'g'$, and "endogamous" or within-class if its two members are from the same class. A configuration exhibits (full) class "stratification" if dyads never cross class boundaries; that is, all dyads are endogamous in that both members are always from the same class hg. This Lemma precisely pins down the payoffs of all stable equilibria under stratification.

Lemma 2. *Let $\mathcal{C}$ be an equilibrium configuration. If j, j' are in a dyad and $h_j = h_{j'}$, $g_j = g_{j'}$, (henceforth if j, j' is an endogamous pair) then $w_j = \frac{\pi(j,\ddot{j})}{2}$, which we denote by $w_d(j)$, and $w_{j'} = w_d(j')$. If instead j is in a monad, $w_j = w_m(j) = \pi(j) = [\alpha + h_j z_j]p - c$.*

This simple payoff characterization is possible because our k-tuples assumption rules out partner scarcity problems under endogamy. We denote the number of journalists in each class by J_{hg} for each of the four classes. By our $k - tuple$ assumption, this number is even so long as journalists of a given type have the same matching behavior, as they must do under stratification. So under stratification, we do not need to worry about scarcity of partners for any given class. We show that in any stable equilibrium, j does not care with which particular partner from j's preferred matching class, j gets matched. So the even k-tuplicates assumption precludes partner scarcity problems. However, with cross-class pairing, or exogamy, the scarce side of the match will get a rent; in addition, rather than random rationing, the journalists on the long side of the "matching market" that gain most from being paired are paired in Lexicographic order and the cut-off type determines the size of the rent taken by the short-side; each member of the scarce class, gets this same rent on top of any marginal return based on its own type. We demonstrate how this works in the example below.

We now provide an expression for the incentives of journalists in arbitrary classes to form a dyad relative to the always available option of creating monads. Distinguishing the values of h, g, z, w, d of j, j' by adding a prime for j', the "pairing gain", PG, from being monads ($d = d' = 0$) is given by,

$$PG_{jj'}(d = d' = 0) = \left[\big((g + g')\alpha + gz' + g'z\big)\sigma + (\sigma - 1)(2\alpha + hz + h'z')\right]p - 2\Delta$$

The last term represents the total increase in the organizational cost. The terms involving g, g' capture the *spillover effects* associated with good neighbors. Notice the complementarity between having high quality journalists (high z) pair up with good neighbors ($g = 1$), since those journalists better exploit the inflow of traffic from good neighbors; e.g., the largest possible spillover is $\sigma(\alpha + 1)p$. The terms involving h, h' capture the *prominence effect* associated with the increased primary traffic attracted by dyads relative to monads: $p_d - p_m = (\sigma - 1)p$. Notice that the spillover effect is always augmented by the prominence parameter σ, since spillovers can only arise with dyads (in the baseline setting).This pairing gain PG is always strictly increasing in g, g', weakly increasing in z, z', and strictly increasing in h, h' provided $\sigma > 1$.[12]

Adding these spillover and prominence effects for a pair with $z = z' = 1, g = g' = h = h' = 1$ reveals the maximum benefit from forming a dyad, $2(\alpha + 1)(2\sigma - 1)p$. To ensure that dyads can potentially arise, we thus need to assume that $\Delta/p \leq (\alpha + 1)(2\sigma - 1)$ for some possible p that is consistent with the existence of dyads.[13] In the case with $\sigma = 1$, p is always equal to $1/J$ and this reduces to assuming $\Delta J \leq \alpha + 1$. To deal with the general case, we suppose that Δ is low enough to ensure a dyad of journalists with $z = g = 1$ is always productive,

Assumption A2: $\Delta/\overline{p} \leq (\alpha + 1)(2\sigma - 1)$ *where* $\overline{p} = \max_{D \geq 1}(J + 2(\sigma(D) - 1)D)^{-1}$.

We now proceed to discuss in detail a family of stable configurations that have in common that dyads are all formed by journalists in class $(1, 1)$. If all dyads are in the same class $(1, 1)$, we know from above that any j in this class has a payoff $w_d(j)$ and all other journalists gets $w_m(j)$. Moreover, the individual rationality implies that the quality of each journalist in a dyad must exceed a cutoff, which since $(h, g) = (1, 1)$, we denote by:

[12]The last claim is true because a rise in h, h' corresponds to a rise in z, z' and the corresponding quality z, z' is necessarily weakly above z_0 which is strictly positive.

[13]For instance, in the case in which $\sigma > 1$ the largest p, and the one under which the inequality is most likely to hold, corresponds to $D = 0$. So although such a value of p satisfies the inequality, it is not consistent with the existence of dyads.

$$z_{11}(D) = \frac{\Delta}{p(2\sigma - 1)} - \alpha$$

Note that the cutoff depends on D through p and possibly through σ and we stress this dependence as it is important for expressing clearly the structure of stable configurations. Importantly, the number of dyads that actually form under any realization of types must be consistent with the threshold $z_{11}(D)$. The following lemma, which can be readily established by treating D as a continuous variable, is the basis for the existence of the simple equilibria described in Proposition 2.

Lemma 3. *$z_{11}(D)$ is strictly increasing in D unless $\sigma = 1$ in which case it is constant.*

Since j must lie in the class $(1, 1)$, it must be the case that $z \geq z_0$, so we define $\hat{z}_{11}(D) = \max\{z_0, z_{11}(D)\}$. Any j with $z \geq \hat{z}_{11}(D)$ *could* form a dyad if users were optimistic enough about its pairing behavior for $g(z) = 1$, but we cannot fully pin down a unique configuration, because many beliefs about $g(\cdot)$ can arise. There are multiple equilibria.[14]

The considerations above lead to the intuitively appealing stable configuration set forth by Proposition 1. We denote by $D^* \in \{0, 1, ..., J/2\}$ the least value of D with the property that there are exactly D eligible journalists of quality at least $\hat{z}_{11}(D)$, that is, $\{j : z_j \geq \hat{z}_{11}(D)\} = D$. We defer the discussion of the existence of such a value of D, until after Proposition 1, where we introduce additional definitions required for describing a more general class of equilibria. As the stable configuration that we are about to describe is a special case of the larger family introduced in Proposition 2, we omit the formal proof that it in fact is stable.

Proposition 1. *If assumptions A1 and A2 hold then the configuration in which for each possible realization of types, the D^* journalists of quality at least $\hat{z}_{11}(D^*)$ form dyads (they may pair up in any way) and journalists of all other qualities form monads is stable.*

There are many other stable configurations. A family of such stable configurations closely related to the one just discussed involves restricting the set of journalists that may form dyads to those with qualities in some arbitrary subset E. Notice that if all low quality journalists are bad neighbors, then by Lemma 3 the only types of journalists that could ever be part of a dyad are those of quality above $\hat{z}_{11}(0)$ (in light of Lemma 3, the lowest possible value of the $\hat{z}_{11}(D)$ threshold). We refer to any subset E, of types weakly above $\hat{z}_{11}(0)$ as a set of eligible types. Given a set of eligible types E, a vector of

[14] We cannot even pin down that $g(z) = 0$ for all $z < z_{11}$, since such journalists never form dyads, but this does not affect relevant outcomes. (One might impose that users never expect a journalist of a quality that never gains from forming a dyad to do so, so they put no weight on this quality being matched to a journalist in a dyad, so long as some other journalist might gain from pairing with the journalist that the user first encounters.)

realized types and any particular type z, we denote by $N(z,E)$ the number of journalists of type at least z. We denote by $D^*(E) \in \{0,1,...,J/2\}$ the least value of D with the property that there are exactly D eligible journalists of quality at least $\hat{z}_{11}(D)$, that is $\{j : z_j \geq \hat{z}_{11}(D)\} \cap \{j : z_j \in E\} = D$. In what follows and for ease of exposition we proceed assuming that J is large and in our exposition ignore the low probability with which the vector of realized types is one for which it does not exist.[15]

We need to ensure that being considered a bad neighbor is enough to dissuade those journalists that are not in E from forming pairs among themselves and from threatening the dyads formed by journalists in E. We call the conditions on the parameters required to achieve this, the exclusion restrictions associated to E, $ER(E)$. Letting $\overline{z}_{E^c}$ denote the highest type not in E and $\underline{z}_E$ the lowest type in E they are given by:[16]

Exclusion Restrictions associated to the set of eligible types E, $(ER(E))$:

(1) *No pairing with other excluded agents:*

$\overline{z}_{E^c} \leq \frac{\Delta}{\overline{p}(\sigma-1)} - \alpha$

(2) *No pairing with high quality, good neighbor monads:*

For all D, $\overline{z}_{E^c} \leq \frac{\Delta}{p(2\sigma-1)} - \frac{\hat{z}_{11}(D^*(E))(\sigma-1)}{(2\sigma-1)} - \frac{\alpha(3\sigma-2)}{(2\sigma-1)}$

(3)*No threats to existing pairs:*

$\overline{z}_{E^c} \leq \frac{\Delta}{p(2\sigma-1)} - \frac{(\sigma-1)}{(2\sigma-1)} + \frac{\underline{z}_E+\sigma}{(2\sigma-1)}$

It is worth emphasizing that if E does not exclude any journalist, that is, when $E = [\hat{z}_{11}(0), 1]$ these conditions are not required. It should also be noted, that these conditions are independent. For some parameter one of them may be redundant but nor for others.

Proposition 2. *Let E be any set of eligible types. If assumptions A1 and A2 hold and the exclusion restrictions associated to E, $ER(E)$, are met then the configuration in which for each possible realization of types, the $D^*(E)$ journalists in E of quality at*

[15] Due to the fact that in our setting there is a discrete number of journalists $D^*(E)$ may not exist. It is worth noting that the probability that this is the case tends to 0 with J. To fix ideas, consider the case in which J is constant and assume that the types are uniformly distributed in the interval $[0,1]$. There always exists D^- such that $\{j : z_j \geq \hat{z}_{11}(D^-)\} \cap \{j : z_j \in E\} > D^-$ and $N(\hat{z}_{11}(D^- + k)) \cap \{j : z_j \in E\} < D^- + k$ The nonexistence of D^* occurs when there is a type that lies precisely in the interval $[\hat{z}_{11}(D^-), \hat{z}_{11}(D^- + k)]$. Note however that the interval $[\hat{z}_{11}(0), \hat{z}_{11}(J)] = [\frac{\Delta}{\sigma J(2\sigma-1)} - \alpha, \frac{\Delta}{J(2\sigma-1)} - \alpha]$ and thus the expected number of types that are realized within that interval is constant and equal to $\frac{\Delta}{k(2\sigma-1)}(1 - 1/\sigma)$. There are however $J/2$ subdivisions of this interval by the $\hat{z}_{11}(D)$ thresholds as D varies from 0 to $J/2$. The expected fraction of $[\hat{z}_{11}(D), \hat{z}_{11}(D+k)]$ intervals that host a type thus tends to 0 as j grows, and this provides an upper bound for the probability that $D(E)$ does not exist. Notice that there is also a direct solution to our problem of nonexistence which specifies that a small fraction of the journalists use mixed strategies.

[16]Note that the only reason why a monadic journalist above $\hat{z}_{11}(D^*(E))$ might be an unattractive partner in a dyad is that she is not a good neighbor. It is therefore the case that the parametric conditions under which we are able to construct equilibria involving agents above quality $\hat{z}_{11}(D^*(E))$ not forming dyads become very stringent if E is very different from $[\hat{z}_{11}(0), 1]$, in the sense that it excludes some very high quality types.

least $\hat{z}_{11}(D^*(E))$ form dyads (they may pair up in any way) and journalists of all other qualities form monads is stable.

The vectors of payoffs that support these configurations are determined by Lemma 2 as $w_m(j)$ and $w_d(j)$ since all pairing is endogamous – only high quality journalists pair up so they are all good neighbors.

The main ideas behind the proof of Proposition 2 are discussed in the preceding paragraphs. In the formal proof of this result (which we provide in the appendix), we must also demonstrate that the pairing gain for all mixed-class pairs are strictly negative, but this is immediate given the lower bound on Δ/p, as derived above.

Furthermore, we can strengthen our results considerably by restricting attention to contexts where the relative cost Δ of forming a dyad instead of a monad is so high that journalists with h or g below 1 never form a dyad. The precise cut-off on Δ is given by considering the gain from a deviation to monads by any paired journalists j, j' where at least one of them has $(h, g) \neq (1, 1)$. Note that the separation gain for j, j' upon forming monads relative to their payoffs as dyads equals minus one times their pairing gain (PG). If any of h, h', g, g' is zero, the bracketed term in $PG_{jj'}$ is bounded above by $(4\alpha + 3)\sigma - (2\alpha + 1)$.[17] So letting $\delta = (\alpha + 1)(2\sigma - 1)$, if we assume $\Delta/\bar{p} > \delta - \frac{\sigma-1}{2}$ and $\Delta/\underline{p} \leq \delta$, all dyads must have both journalists in the class with $(h, g) = (1, 1)$. This region is certainly non-empty if $\bar{p} - \underline{p}$ is small enough, which for instance is the case under fixed σ when σ is closed enough to 1.

Proposition 3. *If Assumptions (A1) and (A2) hold, and additionally, $\Delta/\bar{p} > (\alpha + 1)(2\sigma - 1) - \frac{\sigma-1}{2}$, then:*

(1) All stable configurations are of the form described in Proposition 2.

(2) If we refine the off-equilibrium beliefs of the users by relying on the the intuitive criterion, Cho and Kreps (1987), a configuration is stable if and only if it prescribes a $\boldsymbol{TP}(\hat{z}_{11})$ pairing for every possible realization of the vector of types. That is, if and only if the whole set of types who could ever be part of dyad is eligible ($E = \hat{z}_{11}(0)$).

To see why part (2) of Proposition 3 holds consider a configuration under which $E \neq [\hat{z}_{11}(0), 1]$. As discussed above, the stability of such a configuration entails that $g = 0$ for journalists of qualities not in E. That is, if a user were to reach a journalist of quality in E^c then he must believe that her partner is of low quality. However, in the

[17] This comes from considering the upper bound with h or $h' = 0$, which is certainly below this value (for a tighter bound, one could replace 3 by $2 + z_0$), and noting that the bracketed term is bounded above by $(3\alpha + 3)\sigma - 2(\alpha + 1) = (3\sigma - 2)(\alpha + 1)$ if g or $g' = 0$. We can compare benefits for hypothetical h, g even though $h = 0$ actually requires $z < z_0$. In this hypothetical, being a good neighbor is more important than being a high type: a high type better exploits prominence than a low type but raises the opportunity cost of ceasing to produce as a monad, whereas a good neighbor's spillover always benefits from prominence, adding the fixed revenue α for its pair, since monads never exploit spillovers.

parameter range under consideration, a low quality journalist could not profit from such a deviation (forming an unexpected dyad with a journalist in E^c), even if the user read her article presuming her to be of high quality after reading the article but the unexpected dyadist. That is, no journalist other than a high quality one would profit from forming a dyad, regardless of the beliefs of the user. Thus the only beliefs consistent with the intuitive criterion entail $g = 1$ for all high quality journalists (regardless of whether they are part of a dyad or not). However, if this the case, then all journalists of quality in $[\hat{z}_{11}(D^*), 1]$ must form dyads.

In addition, without the lower bound condition, we can define cutoffs analogous to z_{11}, z_{10}, and z_{10} and construct other stable configurations. In part (2) of Proposition 3 , beliefs about a partner's quality are always increasing in the primary journalist's quality. We now construct an example to show that there exist stable configurations entailing beliefs $\mathbf{g}$ that decrease monotonically in quality. In such equilibria, high quality journalists are bad neighbors and are therefore unable to rely on each other to create spillovers. On the other hand, low quality journalists are good neighbors and act as conduits for high quality journalists.

Example 1. *Let* $z_e = \frac{2\Delta}{p(2\sigma-1)} - \alpha\frac{3\sigma-2}{2\sigma-1}$ *and assume that* $z_0 < z_e \leq 1$ *(Condition 1) and* $z_0 \leq \frac{2\Delta}{p\sigma} + \frac{c_m}{p\sigma} - \frac{\alpha(\sigma-1)}{\sigma}$ *(Condition 2)*

Consider the configuration under which given any realization of the vector of types, all journalists of quality weakly greater than z_e *(henceforth excellent journalists) form as many possible number of pairs between low quality agents and themselves starting with the best high quality journalists and going down the list (matching each of them at random with a low quality journalist). All remaining agents form monads, and the journalists on the short side of the market (Low quality journalists or journalists of quality above* z_e*) obtain an appropriate share of the profits of dyads, in order to prevent the unmatched individuals in the long side from destabilizing their partnerships.*

First consider a type realization in which there are more journalists of quality above z_e *than low quality journalists. In this case all low quality journalists form dyads with excellent journalists, giving precedence to the highest quality ones. We let* $\tilde{z}$ *denote the quality of the highest quality journalist that fails form a dyad. All the unmatched excellent journalists form monads as do all Intermediate quality journalists (those of quality weakly greater than* z_0 *but lower than* z_e*). Let* $w^* = w_j = 3\alpha\sigma p + 2\tilde{z}p\sigma - c_d - (p\alpha + p\tilde{z} - c_m) = \alpha p(3\sigma-1) + \tilde{z}p(2\sigma-1) - c_d + c_m$ *be the payoff of low quality journalists (they are in dyads),* $w_j = p\sigma(3\alpha + 2z_j) - c_d - w^*$ *for high quality members of dyad and* $w_j = p\alpha + pz_j - c_m$ *for monads.*

Now consider a type realization in which there are more low quality journalists than excellent journalists. In this case all excellent journalists form dyads with low quality journalists. Intermediate quality journalists (those of quality weakly greater than z_0 *but* *lower than z_e) create monads, as do all the low quality journalists that do not form a partnership with an excelent journalist. Let $w_j = p(\alpha + h_j z_j) - c_m$ for monads and let $w_j = p\sigma(3\alpha + 2z_j) - c_d - (p\alpha - c_m)$ for high quality members of dyads and $w_j = p\alpha - c_m$, for low quality members of dyads.*

The configuration above determines g for excelent quality journalists and for low quality journalists. In particular, $g_j = 1$ if $z_j < z_0$ and $g_j = 0$ if $z_j \geq z_e$ Let $g_j = 0$ for all intermediate quality journalists $z_0 \leq z_j < z_e$ (this is the only group that is completely outside of the market in equilibrium).

In general, competition between journalists to pair with the most attractive partners allows us to tightly pin down the nature of revenue and cost sharing in dyads. The exact payoff that j receives in a given equilibrium configuration depends on the vector or realized types $\mathbf{z}$, as well as j's own type z, which affects her revenues when her article is read, and which determine her class (h, g). The dependence on the other journalists's types partly occurs via primary traffic p, and partly via the feasibility of forming dyads. We say that a dyad is "exogamous" or mixed class if its two journalists are from different classes $hg \neq h'g'$, and endogamous or within-class if its two members are from the same class. We will say that a configuration exhibits (full) class stratification if dyads never cross class boundaries; that is, both members are always from the same class hg. Let the number of journalists in each class be denoted by J_{hg} for each of the four classes. Note that, given our $k - tuple$ assumption,this number is even so long as journalists of a given type have the same matching strategy, as they do in equilibrium. So under stratification, we do not need to worry about scarcity of a given class: we show that in any stable equilibrium, j does not care with which particular partner from j's preferred matching class, j gets matched, and given the duplicates assumption, there is never a problem of partner scarcity. However, with inter-class pairing, or exogamy, the scarce side of the match will get a rent; in addition, rather than random rationing, the journalists on the long side of the "matching market" that gain most from being paired are paired in Lexicographic order and the cut-off type determines the size of the rent taken by the short-side; each member of the scarce class, gets this same rent on top of any marginal return based on its own type.

Social value of dyads

The following results establish that under plausible differences between the offline and online environments, the online $\mathbf{TP}(\hat{z})$ stable configurations discussed above may be less efficient than the corresponding offline $\mathbf{TP}(\hat{z})$ stable configurations.

Proposition 4. *Consider the family of **TP($\hat{z}$)** configurations. The expected utility of the user and total welfare are strictly decreasing in $\hat{z}$ for all $\hat{z} \geq z_0$ as long as σ is close enough to 1.*

Proof of Proposition 4:

Consider a $\mathbf{TP}(\hat{z})$ configuration and let $F(z_1, z_2)$ denote the associated CDF of the quality of the two articles in dyads.[18] The expected utility conditional on arriving in a dyad is thus:

$$\int\limits_{z_1,z_2} (z_1 + z_2) dF(z_1, z_2) = \int\limits_{z_1} z_1 dF_1(z_1) + \int\limits_{z_2} z_1 dF_2(z_1)$$

where F_1 and F_2 represent the CDFs of the marginal distributions of z_1 and z_2. If the search process does not distinguish between dyads nor among the articles therein contained then F_1 and F_2 are just equal to the distribution $F_{\hat{z}}$ of z conditional on $z \geq \hat{z}$.[19] In what follows we let $q(\hat{z})$ denote the probability that a user is routed by the search process to a dyad, $E_{\hat{z}}$ denote $E[z|z \geq \hat{z}]$ and $E_{z_0,\hat{z}}$ denote $E[z|z_0 \leq z \leq \hat{z}]$. Furtehermore expressions of the form $pr(z \geq z_0|)$ denote the conditional probabilities of events as measured by the type distribution with CDF F and pdf f. The expected utility $E(U(\hat{z}))$ of a user can be written as:[20]

$$\begin{aligned}
& q(\hat{z})2E_{\hat{z}} + (1 - q(\hat{z})) \left(pr(z \geq z_0|z < \hat{z})(z_0 + E_{z_0,\hat{z}}) + pr(z < z_0|z < \hat{z})2z_0\right) \\
= \; & q(\hat{z})2 \int_{\hat{z}}^{1} \frac{zf(z)}{1 - F(\hat{z})} dz \\
+ \; & (1 - q(\hat{z})) \left(\frac{F(\hat{z}) - F(z_0)}{F(\hat{z})} \left(z_0 + \int_{z_0}^{\hat{z}} \frac{zf(z)}{F(\hat{z}) - F(z_0)} dz \right) + 2z_0 \frac{F(z_0)}{F(\hat{z})} \right)
\end{aligned}$$

We can thus compute $\frac{\partial E(U(\hat{z}))}{\partial \hat{z}}$ as:

$$\begin{aligned}
& q(\hat{z}) \left(\frac{-2\hat{z}f(\hat{z})}{1 - F(\hat{z})} + \frac{2E_{\hat{z}}f(\hat{z})}{1 - F(\hat{z})} \right) + \frac{\partial q(\hat{z})}{\partial \hat{z}} 2E_{\hat{z}} \\
+ \; & (1 - q(\hat{z})) \left(\frac{\hat{z}f(\hat{z})}{F(\hat{z}) - F(z_0)} - \frac{f(\hat{z})}{F(\hat{z}) - F(z_0)} E_{z_0,\hat{z}} \right) \frac{F(\hat{z}) - F(z_0)}{F(\hat{z})} \\
+ \; & (1 - q(\hat{z})) \left(\frac{f(\hat{z})F(z_0)}{F(\hat{z})^2} (z_0 + E_{z_0,\hat{z}}) - \frac{2z_0F(z_0)f(\hat{z})}{F(\hat{z})^2} \right) \\
- \; & \frac{\partial q(\hat{z})}{\partial \hat{z}} \left((z_0 + E_{z_0,\hat{z}}) \left(1 - \frac{F(z_0)}{F(\hat{z})} \right) + 2z_0 \frac{F(z_0)}{F(\hat{z})} \right)
\end{aligned}$$

Some grouping and rearrangement of terms leads to:

[18] Notice that the CDF incorporates the initial distribution of the types, any details of the configuration, and of the random search process leading users to their first article. Our assumption is that conditional on arriving at a dyad selected by the search process and given any for any $z_1, z_2 \in [\hat{z}, 1]$ the probability that the qualities of the first and second articles encountered is at most z_1 and z_2 is well defined and given by $F(Z_1, Z_2)$.

[19] And importantly, it is independent of the number of realized dyads.

[20] For convenience, we do not denote the dependence of the user's expected utility on the other parameters.

$$
\begin{aligned}
& q(\hat{z})\frac{2f(\hat{z})(E_{\hat{z}}-\hat{z})}{1-F(\hat{z})}+(1-q(\hat{z}))\frac{f(\hat{z})}{F(\hat{z})}\left((\hat{z}-E_{z_0,\hat{z}})+\frac{F(z_0)}{F(\hat{z})}(E_{z_0,\hat{z}}-z_0)\right) \\
+ \quad & \frac{\partial q(\hat{z})}{\partial \hat{z}}\left(\frac{F(z_0)}{F(\hat{z})}(E_{z_0,\hat{z}}-z_0)+(E_{\hat{z}}-E_{z_0,\hat{z}})+(E_{\hat{z}}-z_0)\right)
\end{aligned}
$$

Notice that given that $(E_{\hat{z}}-\hat{z})$, $(\hat{z}-E_{z_0,\hat{z}})$, $(E_{z_0,\hat{z}}-z_0)$, $(E_{\hat{z}}-E_{z_0,\hat{z}})$ and $(E_{\hat{z}}-z_0)$ are all nonnegative for any type distribution we conclude that $\frac{\partial E(U(\hat{z}))}{\partial \hat{z}}<0$ if and only if:

$$
\begin{aligned}
& q(\hat{z})\frac{2f(\hat{z})(E_{\hat{z}}-\hat{z})}{1-F(\hat{z})}+(1-q(\hat{z}))\frac{f(\hat{z})}{F(\hat{z})}\left((\hat{z}-E_{z_0,\hat{z}})+\frac{F(z_0)}{F(\hat{z})}(E_{z_0,\hat{z}}-z_0)\right) \\
< \quad & -\frac{\partial q(\hat{z})}{\partial \hat{z}}\left(\frac{F(z_0)}{F(\hat{z})}(E_{z_0,\hat{z}}-z_0)+(E_{\hat{z}}-E_{z_0,\hat{z}})+(E_{\hat{z}}-z_0)\right)
\end{aligned}
$$

In particular it can be seen that if the prominence function is such that $\frac{\partial q(\hat{z})}{\partial \hat{z}}\geq 0$ then certainly $\frac{\partial E(U(\hat{z}))}{\partial \hat{z}}>0$. More generally however $\frac{\partial q(\hat{z})}{\partial \hat{z}}$ will tend to be negative and the overall effect depends on the details of the process.

In particular when $\sigma=1$, then $q(\hat{z})$ is simply $(1-F(\hat{z}))$ and $\frac{\partial q(\hat{z})}{\partial \hat{z}}=-f(\hat{z})$. Replacing in the expressions above, we obtain $\frac{\partial E(U(\hat{z}))}{\partial \hat{z}}=f(\hat{z})(z_0-\hat{z})$ which is negative as long as $\hat{z}>z_0$, as is the case in the family of equilibria considered in Proposition 2.

We draw these ideas together in the following proposition, which supposes that the online world differs from the offline world in reducing the fixed cost from c to $c-C$ and no other change. It follows that c_m diminishes by C while $\frac{c_d}{2}$ diminishes by only $C/2$. The per capita added cost of forming a dyad, Δ, therefore rises by $C/2$, reflecting the point that with a low fixed cost reduces the net gain from forming a dyad because the fixed cost saving falls. Now this causes z_d to rise by $\frac{C/2}{p(2\sigma-1)}$. So that in the focal equilibrium (from part 3 of Proposition 2), the degree of pairing into dyads falls and consumers are worse off. (The result can be generalized but it is clearest where the equilibria are unique for all comparative statics.)

Proposition 5. *Under the assumptions of part 2 of Proposition 3, the shift to lower fixed costs associated with online technology reduces the number and fraction of journalists who pair up. By Proposition 4, for σ close to 1, this reduces users' expected utility and total welfare.*

4 Discussion and motivation of assumptions

4.1 Remarks on prominence and signals

Our main motivation behind the minimalist discriminating role of prominence is simply to focus all attention on the role of the peer interactions between journalists. By making the task of finding high quality content more difficult, this rigid prominence function lets us home in on how the decentralized decisions of a distributed set of individual journalists can convey information about quality to users. In the case of $\sigma = 1$, each user's primary article is equally likely to be from any of the journalists in the set of potentially relevant journalists, J, fully independent of their qualities $\mathbf{z}$. The online world's great surfeit of content, combined with taste heterogeneity, unavoidably generates large equivalence classes of material that is essentially undifferentiated by content rankings. This equivalence class may include material from organizations of different sizes. We effectively focus on how users find content from with this relevant set J of equivalent material.

Nonetheless, we briefly mention how one can motivate the minimalist approach in two specific cases: online platforms that use algorithms to make recommendations and recommendations from close social contacts that do not rely on digital technologies.

For the interpretation based on search engines and the algorithms used by online social networks, it is natural to suppose that they can detect dyads and monads. While it is extreme to suppose that they learn nothing about quality, this captures in a simplified way the idea that they learn about quality from user behavior and peers' linking choices, of which user-learning necessarily involves delays and the peer-linking effect is here represented by σ.

In the baseline, peers' linking choices reduce to their dyad/monad status d_j and is endogenously informative of quality as we show below. Learning from how long users spend on different websites is valuable, especially for walled gardens like Facebook that can fully monitor user behavior within Facebook, but clearly trivial when journalists are totally new since there is no past user behavior to learn from.[21] In this extreme case, the role of the search engine or news aggregator may be simply to recommend articles in topic areas that may interest the user, or to exclude an unmodeled fringe of "journalists" so lousy, they do not merit the name of journalist. Heterogeneity in topics will be a key extension for the analysis, especially as search engines are well suited for this and provide an advantage of the online environment from which we currently abstract.

Notice that in this search engine interpretation, we implicitly suppose the user inspects a suggested result and has a cost of doing so that makes it prohibitive to inspect more than one result. After all, it would be implausible to imagine the search engine only finds and offers one result. The search engine itself does not observe the quality signal

[21] In addition, platforms run by entities such as Google and Facebook may introduce their own biases; see Burguet et al. (2015).

that is learned by the user on inspecting the engine's suggested links. This restriction on the number of searches may reflect the user's time constraints or distaste for having to adopt a critical frame of mind, as may be needed to inspect the quality of an article; for instance, it may feel more pleasant to simply trust the writer, and trusting may make it difficult to inspect whether the article is good or not. In our scenario, performing a second search would be prohibitively costly so that if a first search fails to deliver a sufficiently high quality, the user strictly prefers to resort to the default outlets, where he knows how to efficiently obtain the material that he is interested in, rather than undertake a second search. We discuss this further when we (later) extend the model to treat the case where users only learn quality by experiencing the article and by then it becomes too late to switch to another article.

For the interpretation where a friend or respected contact of each user recommends the primary article to that user, we can more readily suppose that the recommender has just observed one article, so that it is fine if he observes its quality and tells this to the user (without any user inspection). Naturally, matters change if the recommender reads both articles in a dyad before recommending; we discuss how a similar set of results arise in such settings later on. For now, we merely note that $\sigma = 1$ can be motivated by the scenario in which the recommender reads just one article and explains his experience to his co-user, independent of whether it was in a dyad or monad. The model would also have to change to accomodate recommenders who have multiple experiences and only bother to relate the most useful ones. Such recommender behavior would raise the prominence of better articles and better linked articles. (Even then, if the recommender's motive is purely to discuss a given article together, the value of its links would not play a role, though he would never gain the chance to discuss a sub-z_0 article.)

The constant σ may well be different online and offline. In particular, PageRank is an important ingredient in online search algorithms and this favors dyads over monads, as they have higher centrality; this implies $\sigma > 1$. In addition, $\sigma(D)$ tends to be a decreasing function since PageRank promotes nodes based on relative, rather than absolute, connectedness. Offline "word of mouth" processes may also favor larger groups and the effect is greater if there are few groups with a size advantage. One can endogenize this case later.

While one cannot fully rule out $\sigma < 1$, several factors push towards greater prominence of dyads. Dyads offer more content and this tends to benefit users, so recommenders who want to attract or help users have a good motive to favor dyads. The average quality on dyads can be lower in some equilibria, but equilibria in which better journalists self-select into dyads are more robust. A thorough treatment would need to have σ endogenous and also explicitly model stochastic variation in topic relevance or noisy quality signals that could explain our continuous prominence function.

4.2 Advertising revenues

We do not distinguish between the advertising value generated by primary and secondary traffic, but in the online setting primary visits may occur via homepages distinct from the individual article webpages. We sometimes primary traffic as landing traffic for the case of a (web)site. In the offline world, the interpretation is different. Even in the offline scenario where a user whose initial search strategy recommends a particular article, this person can be said to "land" on that article if they go first to that article as soon as they acquire the newspaper in which it is contained. Of course, in practice, a user may be distracted by the frontpage or contents page; online, there may be a homepage, similar to the frontpage of a legacy paper, but online search may more directly lead to landing on a specific article page with no risk of distraction. We neglect the risk of distraction for now. We do distinguish between homepages and article pages in an extension, but that subtlety is not necessary for our key results.

5 Conclusion

In brief, we have provided conditions under which one can expect greater difficulties for people to find high quality content online than used to arise in the offline world. As well as inducing a less discriminating self-selection of journalists, reducing the entry cost of publishing reduces the impetus for journalists to pair up into larger organizations and this complicates the search for quality content, because grouping up generated a positive informational externality for users. Evidence of high concentration of online news traffic on a small number of websites, especially those of the conglomerates that were dominant before the internet grew strong, is at least consistent with our prediction. Nonetheless, we are at an early stage in the analysis, since we have abstracted from many important additional differences. For instance, the internet environment makes topic search far easier and the costs of switching between outlets is typically much reduced, especially when individual outlets provide hyperlinks to other journalists, as compared to the offline setting.

Appendix

Proof of Lemma 2:
If $j' = \ddot{\jmath}$, the conclusion follows immediately from equal treatment. Now suppose $j' \neq \ddot{\jmath}$. Equal treatment and stability applied to the twin pair $(j, \ddot{\jmath})$ imply: $2w_j = w_j + w_{\ddot{\jmath}} \geq \pi(j, \ddot{\jmath}) = 2w_d(j)$, so $w_j \geq w_d(j)$. Denoting the common value of $h_j, h_{j'}$ by h and of $g_j, g_{j'}$ by g, feasibility implies: $w_j + w_{j'} \leq \pi(j, j') = \pi(j,j)/2 + \pi(j',j')/2 = w_d(j) + w_d(j')$ because,

denoting the common value of $h_j, h_{j'}$ by h and of $g_j, g_{j'}$ by g,

$$\begin{aligned}\pi(j,j') &= \big((2\alpha + h_j z_j + h_{j'} z_{j'}) + (g_{j'}(\alpha + z_j) + g_j(\alpha + z_{j'})\big)p_d - c_d \\ &= \big((2\alpha + h(z_j + z_{j'})) + g(2\alpha + z_j + z_{j'}\big)p_d - c_d\end{aligned}$$

which is linear in $(z_j + z_{j'})$ and therefore equals half the sum of its values for (j,j) and (j',j').

Together, $w_j \geq w_d(j)$, $w_{j'} \geq w_d(j')$ and $w_j + w_{j'} \leq w_d(j) + w_d(j')$ imply $w_j = w_d(j)$ and $w_{j'} = w_d(j')$. ∎

Proof of Proposition 2:

Notice the configurations described in statement of the Proposition only involve dyads among high quality journalists. It follows that for any member j of a dyad it must be the case that $g_j = 1$, and thus all dyads are endogamous pairs: $h_j = h_{m(j)} = 1$ and $g_j = g_{m(j)} = 1$. From Lemma 2 it follows that the payoffs of any journalist j that is part of dyad is $w_j = w_d(j) = \big((\alpha + hz_j) + g(\alpha + z_j)\big)\sigma p - c - \Delta = (2\alpha + 2z_j)\sigma p - c - \Delta$. The payoffs of monadic journalists in turn just equal the profits of their outlets and are thus $w_j = \alpha p - c$ in the case of low quality journalists and $w_j = (\alpha + z_j)p - c$ in the case of high quality ones. Finally, in order to assess the stability of all the pairings that might be induced by the configuration we need to specify the out of equilibrium beliefs that the users rely on in order to decide how to act in case they encounter a journalist of a quality unexpected in a dyad. So we let $g_j = 0$ for all j such that $z_j \notin E$ (where E is the arbitrary set of eligible types), or $z_j < \hat{z}_{11}(0)$. There are four (h,g) classes of journalists whose incentives we must consider $(0,0)$, $(1,0)$ monadic $(1,1)$ and dyadic $(1,1)$.

Monads

We begin by studying the incentives of monads to form pairs among them. Note that the pairing gain for two monadic journalists,

$$PG_{jj'}(d = d' = 0) = \big[\big((g + g')\alpha + gz' + g'z\big)\sigma + (\sigma - 1)(2\alpha + hz + h'z')\big]p - 2\Delta$$

is increasing in their classes (h,g) and (h',g') and in their qualities, so we only need to make sure that (1) a journalist of highest quality of class $(1,0)$, potentially of quality z_E prefers not to deviate with an identical identical journalist (2) that this same journalist will not deviate with the highest quality user of class $(1,1)$, potentially of quality just below $\hat{z}_{11}(D^*)$ (3) that the journalist of class $(1,1)$, of quality just below $\hat{z}_{11}(D^*)$ does not want to pair with an identical journalist. (1) and (2) are guaranteed by the exlusion restrictions associated to E, $ER(E)$. (3) Is guaranteed by the definition of $\hat{z}_{11}(D^*)$.

Now we move on to contemplate the incentives of monadic journalists to pair with journalists that are already part of a dyad. As before, the pairing gain is increasing in z, z' and g.

$$PG_{jj'}(d=0, d'=1) = \big[z(h(\sigma-1)+\sigma) + z'\sigma(g-1) + \alpha(\sigma(g+1)-1)\big]p - \Delta$$

h' and g' play no role as all the dyads found in the configurations in question involve high quality journalists that are good neighbors. We thus just need to check the threats posed by (1)The highest quality not in E (or if E^c is empty, a quality just below $\hat{z}_{11}(0)$) and the highest quality in journalist E and (2) A quality just below $\hat{z}_{11}(D^*)$ and the highest quality in journalist E . (1) Is guaranteed by the third condition in $ER(E)$ and (2) By the definition of $\hat{z}_{11}(D^*)$.

Finally the definition of $\hat{z}_{11}(D^*)$ guarantees that journalists that are part of dyads find it individually rational not to deviate by becoming monads. Moreover, their payoffs $w_d(j)$, imply that they are indifferent between remaining with their partner and associating with anyone else. In particular $W_d(j) + w_d(j') = \pi(j, j')$ for all journalists that are part of a dyad, regardless of whether $m(j) = j'$ or not. ∎

Discussion of Example 1:
In what follows we analyze in detail Example 1 and we do so by breaking it into two: (A) type realizations in which excellent journalists are on the long side of the partnership market, and (B) type realizations in which low quality journalists are on the long side of the partnership market, The case in which low quality journalists are on the short side is simpler, given that there are no monadic high quality journalists posing a threat to existing dyads and we therefore omit it.

(A) Excellent Journalists are on the long side

(A1)Low quality journalists:
It is individually rational for a low quality journalist to form dyads as long the quality of the journalist that defines their payoff z_{ind} is higher than $\frac{2\Delta}{p(2\sigma-1)} - \alpha\frac{3\sigma-2}{2\sigma-1}$ which is precisely z_e and therefore true. For that same reason she would never form a dyad with a good quality journalist of quality below z_{ind}.
Notice that since their payoff is constant they have no incentive to pairwise deviate with any other already partnered excellent quality journalist. Making sure that they do not wish to form a dyad with another low quality journalist is a different story since these are good neighbors. We require $2w^* \geq (2\alpha + z_j + z_{j'})p\sigma - c_d$ for all $z_j, z_{j'} < z_0$. This is equivalent to: $2(3\sigma p + 2z_{ind}\sigma - c_d - (p\alpha + pz_{ind} - c_m)) \geq (2\sigma + z_j + z_{j'})\sigma p - c_d$. A sufficient condition is $z_0 \leq \frac{2\Delta}{p\sigma} + \frac{c_m}{p\sigma} - \frac{\alpha(\sigma-1)}{\sigma}$ (Condition 1).

(A2) High quality journalists that are in Dyads:

It has already been shown that a such a journalist would not pairwise deviate with a low quality journalist. The condition that he does not pairwise deviate with another paired excellent journalist can be shown to be equivalent to $z_{ind} \leq \frac{\Delta}{p(\sigma-1)} - \alpha$ and a sufficient condition for this to be the case is $z_{ind} \leq \frac{\Delta}{p(\sigma-1)} - \alpha$ which is always true given the definition of z_e.

References

Bar-Isaac, H. and Tadelis, S. (2008). *Seller reputation.* Now Publishers Inc.

Becker, G. S. (1973). A theory of marriage: Part I. *Journal of Political Economy*, pages 813–846.

Burguet, R., Caminal, R., and Ellman, M. (2015). In google we trust? *International Journal of Industrial Organization*, 39:44–55.

Cabral, L. and Hortacsu, A. (2010). The dynamics of seller reputation: Evidence from ebay. *The Journal of Industrial Economics*, 58(1):54–78.

Cho, I.-K. and Kreps, D. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–221.

Dellarocas, C., Katona, Z., and Rand, W. (2013). Media, aggregators, and the link economy: Strategic hyperlink formation in content networks. *Management Science*, 59(10):2360–2379.

George, L. M. and Hogendorn, C. (2012). Aggregators, search and the economics of new media institutions. *Information Economics and Policy*, 24(1):40–51.

Ghatak, M. (1999). Group lending, local information and peer selection. *Journal of development Economics*, 60(1):27–50.

Ghatak, M. (2000). Screening by the company you keep: Joint liability lending and the peer selection effect. *The Economic Journal*, 110(465):601–631.

Jackson, M. O. and Wolinsky, A. (1996). A strategic model of social and economic networks. *Journal of Economic Theory*, 71:44–74.

Jeon, D.-S., Esfahani, N. N., et al. (2012). News aggregators and competition among newspapers in the internet. *American Economic Journal: Microeconomics.*

Legros, P. and Newman, A. F. (2002). Monotone matching in perfect and imperfect worlds. *Review of Economic Studies*, 69(4):925–942.

Legros, P. and Newman, A. F. (2007). Beauty is a beast, frog is a prince: assortative matching with nontransferabilities. *Econometrica*, 75(4):1073–1102.

Liu, Q., Mailath, G. J., Postlewaite, A., and Samuelson, L. (2014). Stable matching with incomplete information. *Econometrica*, 82(2):541–587.