

NBER WORKING PAPER SERIES

THE PRODUCTION AND CONSUMPTION OF SOCIAL MEDIA

Apostolos Filippas
John J. Horton
Elliot Lipnowski
Prasanna Parasurama

Working Paper 28666
<http://www.nber.org/papers/w28666>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, MA 02138
April 2021, Revised April 2025

We thank Hunt Allcott, Erik Brynjolfsson, Paul Oyer, Evan Sadler, Jesse Shapiro, Alex Tabarrok, and Johan Ugander for their very helpful comments, and Jie Lu for his excellent research assistance. Author contact information, code, and data are currently or will be available at <https://www.apostolos-filippas.com>. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Apostolos Filippas, John J. Horton, Elliot Lipnowski, and Prasanna Parasurama. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Production and Consumption of Social Media

Apostolos Filippas, John J. Horton, Elliot Lipnowski, and Prasanna Parasurama

NBER Working Paper No. 28666

April 2021, Revised April 2025

JEL No. D51, D71, D85

ABSTRACT

We model social media as collections of users producing and consuming content. Users value consuming content but due to scarce attention they may not value all content from other users. Users also value receiving attention, creating the incentive to attract an audience by producing valuable content but also through attention bartering—users mutually becoming each others' audience. Attention bartering shapes substantially the patterns of production and consumption on social media, explains key features of social media behavior and platform decision-making, and yields sharp predictions that are consistent with data we collect from #EconTwitter. We conduct Twitter and Instagram user surveys that yield additional direct evidence in support of attention bartering, and we discuss the implications of attention bartering for the design of social media platforms.

Apostolos Filippas
Gabelli School of Business
Fordham University
140 W 62nd Street
New York, NY 10023
apostolosfilippas@gmail.com

John J. Horton
Massachusetts Institute of Technology
Sloan School of Management
100 Main St
Cambridge, MA 02142
and NBER
john.joseph.horton@gmail.com

Elliot Lipnowski
Columbia University
e.lipnowski@gmail.com

Prasanna Parasurama
New York University
44 West 4th Street
KMC 8-186
New York, NY 10012
pparasurama@stern.nyu.edu

1 Introduction

Traditional media, such as newspapers, TV, and radio, are characterized by a few “stars” who produce content, and a much larger group of people who purely consume that content (Rosen, 1981; Krueger, 2005). On social media, by contrast, a large number of users both produce and consume content: the “social” of “social media” is that millions of people tweet on Twitter, dance on TikTok, rant on Facebook, open boxes on YouTube, share sunsets on Instagram, and announce exciting new professional chapters on LinkedIn. However, all of this new content has not been matched by an equal increase in the total supply of human attention—the scarce factor consumed during content consumption (Simon, 1971). This paper examines how human attention is allocated in equilibrium over social media content—who consumes what—and how this allocation in turn affects the incentives for content production.

A starting point is to observe that the costs of producing, distributing, and accessing content have fallen dramatically with digitization, and with the proliferation of modern computing technology. Today, marginal distribution costs are de minimis, and technical production costs have diminished radically, even for audio and video content. Falling costs can explain the larger number of professional producers, the greater variety of consumption, and why users spend more time on social media. Nevertheless, producing content still has an opportunity cost; why then do so many amateur users produce content on social media, when almost no one is paid to do it?

Part of the explanation for amateur social media production is the desire for an audience, and a positive reaction from that audience. Although people sometimes produce content without an external audience—diaries exist, and people enjoy playing music by themselves—many clearly value having an audience. The desire for an audience can explain the effort put into social media—i.e., the production decision—but it can also affect the *consumption* decision, in a way that is particular to social media.

We argue that a distinguishing feature of social media platforms is that they allow users to exchange some of their own attention for the attention of others in order to obtain a larger audience. Unlike most economic situations, in which individuals produce as a means to consume, social media users can partially consume in order to produce. In particular, social media users can “follow” others not (only) to consume their content, but to be “followed back” and consume their attention. We call this exchange *attention bartering*.

Throughout this paper, we focus on reciprocal following because it is a convenient proxy for attention bartering—“I will follow you only if you follow me.” But attention bartering goes beyond reciprocal following to reciprocal engagement. On platforms where following relationships are bilateral, such as Facebook or LinkedIn, attention bartering manifests as users reciprocally engaging with each others’ content—“I will engage with your posts only if you engage with my posts.”

Consider Alice, a user who decides where to allocate her attention. On traditional media, Alice can use her attention only to consume content, and hence she allocates her attention toward consuming the best possible content—content produced by “stars” (which may or may not be specific to Alice’s tastes). Alice also likes receiving attention, but her content alone does not attract an audience—in terminology we use throughout, her ability is lower than that of the stars.¹ On social media, however, Alice can strike an implicit “deal” with the similarly good-but-not-great user Bob. Alice agrees to give Bob some attention, in exchange for Bob giving some attention to her. With this deal in place, they both get an audience member, and consequently benefit from producing content. If Bob follows Alice then we say that Bob is Alice’s follower, and that Alice is Bob’s followee.

Both Alice and Bob would have not attracted any followers without the ability to attention barter, and hence they would not have produced any content in a world without attention bartering. The cost of their deal is the attention that they would have preferred to allocate elsewhere: if Alice were to renege on their deal and stop following Bob, Bob would also renege because Alice’s content is not good enough to attract his attention “organically.” Their deal is sustained because social media platforms make it easy to verify that each maintains their end of their bargain. Importantly, because attention bartering requires consuming what the partner produces, both parties have an incentive to seek out the best bargain, that is, the best producers who are willing to barter with them. Together, these incentives create pressure for a kind of assortative matching.

We formalize Alice and Bob’s situation in a model of attention bartering, solving for the equilibrium network structures it generates.² In our model, users have limited attention and are heterogeneous, and attention from other users has positive but diminishing returns. We show that users never attention barter if at least one of them would have organically followed the other in a counterfactual world without bartering. Hence, because it cannot crowd out organic following, the possibility of attention bartering benefits all users—a finding robust to vertical differentiation, horizontal differentiation, and equilibrium selection. The decision of two users to barter depends on whether *both* users have a marginal benefit of an additional follower that exceeds the marginal cost of expending costly attention on the other’s content. Because this marginal benefit depends on a user’s overall “popularity,” a given pair’s optimal bartering decision cannot be assessed in isolation. For this reason, there can be multiple equilibrium networks, with different welfare properties.

We next set aside horizontal differentiation, and focus on a special case of the model

¹We use “ability” here and throughout the paper to refer to the consumption value others derive from a user’s content, without any implied judgement of the objective worth of that content. It is worth noting that consumption value is consumer-specific in the general version of our model, but it is consumer-specific in the stylized “vertical” specialization of our model that guides our empirical application.

²Although we abstract from content production decisions, we interpret our results as though a user produces content if and only if she has followers.

(exhibiting a unique equilibrium) with substantial vertical differentiation among users. As such, users are better off bartering with users who produce better content, all else equal, and two types of users do not attention barter in equilibrium: (1) “stars” with high enough ability to attract followers organically, who therefore have no incentive to barter, and (2) “lurkers” with such low ability that others are not willing to barter with them. Between these two extremes, attention bartering “clubs” emerge, whose “members” have similar abilities, barter with all other members of the same club. Thus, the equilibrium network takes a core-periphery form, with the core followed organically by everybody, and the periphery partitioned into a number of bartering cliques. The size of a club is determined by its members’ incentives—specifically, their willingness to follow its lowest-ability member—and so clubs consisting of lower-ability members must be smaller to generate a higher marginal benefit of attracting a follower. An important consequence of the previous observation is that, among non-stars, users with higher ability in fact follow more other users. We should stress that the core-periphery network, and the club structure comprising the network’s periphery, are not assumed explicitly, but rather emerge in the equilibrium of the vertical model due to the production and consumption incentives of social media users.

Users in bartering relationships—club members—would not have attracted any followers, and would not have produced any content without attention bartering. With attention bartering, club members attract some followers and produce content actively, thereby experiencing lower consumption utilities, but higher attention utilities and higher total utilities from using the platform. As such, attention bartering results in more active users, and more content produced on the platform, albeit of lower average quality.

To assess some of the vertical model’s assumptions and predictions empirically, we collect data from #EconTwitter, a Twitter community comprising professional and amateur users who tweet mostly about economics, and often follow each other. An important feature of our data is that it allows us to observe the production and consumption decisions of #EconTwitter users, as well as to obtain measures of attention bartering and user ability. Note, our broader purpose is to study the formation of user networks on social media platforms—an industry commanding more than \$150bn in annual revenue, and serving as an important information source for its users.³ Although #EconTwitter constitutes a small segment of the population on only one such platform, it serves as a convenient and familiar testing ground to assess some observable predictions of our model.

The distributions of #EconTwitter users’ follower and followee statistics are consistent with our key assumptions of (a) vertical differentiation in ability, and (b) scarce attention.

³For an estimate of social media advertising spending in 2021, see <https://www.statista.com/study/36294/digital-advertising-report-social-media-advertising>. A recent PEW Research Center study estimates that about 11% of US adults prefer to get their news through social media; for more details, see <https://www.pewresearch.org/fact-tank/2021/01/12/more-than-eight-in-ten-americans-get-news-from-digital-devices>.

Follower counts are highly right-skewed, with some “stars” attracting enormous numbers of followers—consistent with producers facing no distribution costs and being vertically differentiated, with the “best” users having vast audiences. In contrast, the distribution of users’ followees is substantially less right-skewed, and does not have a long tail. This feature is consistent with consumers having limited attention: there are no “super consumers” that willingly consume 100 or 1000 times the content consumed by the typical user. Despite this follower/followee distinction—and some star users having enormous ratios—the overall distribution of follower-to-followee ratios is clustered around one. These patterns are consistent with findings reported in previous work spanning several social media platforms, and over a long period of time (Kwak, Lee, Park and Moon, 2010; Myers, Sharma, Gupta and Lin, 2014; Sadri, Hasan, Ukkusuri and Lopez, 2018). Many models of network formation can be made consistent with the patterns described above; an important distinguishing feature of our model is its predictions about outcomes and behaviors conditional upon user ability.

We cannot observe directly user ability, and attempting to infer it from network statistics might make the empirical relationships we find between inferred ability and other network features tautological.⁴ To circumvent this issue in selecting a proxy for ability, we use the number of appearances a user makes on Twitter “lists” curated by other users, adjusted for the tenure of each user at the time of our data collection. Lists allow users to consume the content of a listed member without affecting their follower and followee count. Importantly, because lists are not salient within users’ profiles, prospective followers cannot use list memberships to infer ability. Because of these features of lists, our working hypothesis is that list memberships better capture the true consumption value users derive from a given user’s tweets, being untainted by attention bartering considerations.

Consistent with our ability measure being a good proxy for what users want to consume, the average number of followers is strongly increasing in their ability. In contrast to followers, the average number of followees is first increasing in ability, attains its maximum at medium-to-high ability levels, and then sharply decreases at the highest ability levels. This “followee dip” is a hallmark of attention bartering in our model—the most able users attract many organic followers, having little need to barter, and so they follow only accounts they enjoy following. In our data, we observe who follows whom, and we use reciprocal following as a proxy for attention bartering. We find that low-ability users form reciprocal following relationships with nearly any willing reciprocator, i.e., any user who is willing to “follow back.” As users’ abilities increase they become more selective, following reciprocally higher-ability users on average. Among all users, medium-ability users engage in reciprocal following most

⁴Even so, our analysis still generates nontrivial testable predictions that do not concern ability. For example, in the equilibrium of the vertical model, a user’s number of followees is quasiconcave in her number of followers. However, including both network features and a proxy for ability yields more stringent empirical tests for the model.

intensely. In stark contrast to reciprocal following, no such pattern exists in unreciprocated following: the content that users are willing to consume organically is unrelated to their own abilities, and is of high average quality. These patterns are all predicted by our vertical model of attention bartering; other popular models of network formation do not predict them.

To supplement our analyses and provide additional direct evidence of reciprocity on Twitter, we conducted a survey of Twitter users to directly ask them about their following, unfollowing, and engagement behavior. In the aggregate, users report that they are more likely to follow other users who follow them, more likely to unfollow others who unfollow them, and more likely to engage with others who engage with them. We replicate our user survey with a distinct sample of Instagram users, and we find strong evidence that attention bartering generalizes to other social media. These patterns of reciprocal behavior are consistent with the results of an empirical exercise where we show reciprocity is the strongest predictor of relationship formation on our Twitter data.

While our chief purpose is to offer a novel perspective on social media, the economic phenomena we highlight have social and economic antecedents. For example, consider an “open mic night” at a cafe. The participants are not paid for the content they produce, but rather barter for attention: would-be poets listen to the poetry of their fellow poets, in exchange for having an audience for their own poetry. The cafe owner—like the operators of social media platforms—provides the meeting place for this exchange, profiting from the sale of food and drinks. They can happily sit silently in the background, as long as the poets do not get too rowdy. In the context of academia, paper-reading seminars among graduate students, as well as all-day workshops, exhibit similar economics: academics listen to others’ work in exchange for having others listen to their own work. Similarly, presenting one’s work at a workshop without attending others’ talks is a faux pas, safely available only to the most famous and high-ability academics. For both the open mic night and the workshop, individuals would like to be members of the best club that will have them, as they would prefer to hear better-written poems, and to attend presentations of higher-quality research. However, as in our model, the higher the quality of the other members the more difficult club membership is to obtain.

The platform assumes a passive role in our model but attention bartering can explain several social media strategy decisions. Platforms have an incentive to encourage attention bartering because, as we show, it increases the surplus of users who barter. Through this channel, attention bartering induces users to be more active on the platform, and in turn creates more opportunities to monetize this activity, such as through showing advertisements. Platforms can encourage attention bartering in several ways including rendering follower and followee content highly salient; creating reminders of the presence of the audience (such as “likes”); allowing users to verify easily whether they are being followed by others; and instituting rules against aggressive follow-churn behavior (a kind of defection that undermines

attention bartering). Algorithmic curation of the feed spares the user from some of the disutility of following low-ability users, but without giving up the attention of having that low ability user as a follower. However, the platform must walk this line carefully because algorithmic curation weakens the promise that any follower is in fact seeing the followed user’s content, weakening the attention utility that motivates production and bartering.

This paper joins a burgeoning literature on the economics of social media (Bakshy, Messing and Adamic, 2015; Allcott and Gentzkow, 2017; Alatas, Chandrasekhar, Mobius, Olken and Paladines, 2019; Allcott, Braghieri, Eichmeyer and Gentzkow, 2020; Berman and Katona, 2020; Levy, 2020). Our work focuses on the importance of users’ scarce attention in both their production and consumption decisions.⁵ Our paper is most closely related to Srinivasan (2023), who shows that social media users value attention in and of itself, and receiving attention increases their content production. A key differentiating aspect of our paper is that we model the effects of scarce attention on both the production *and* consumption decisions on social media to explain the formation of social media networks. As such, our model can be viewed as an economic network formation model, with rational economic agents creating and severing links to maximize their utilities. In particular, our model is similar in spirit to the pairwise stability notion of Jackson and Wolinsky (1996).

The rest of the paper is organized as follows. Section 2 provides a brief introduction to Twitter as a paradigmatic social medium that allows for attention bartering. Section 3 presents a model of attention bartering, and derives general results about its equilibrium properties and welfare implications. Section 4 develops a special case of the model (that informs our empirical analysis) in which users are vertical differentiated, and derives its unique equilibrium network. Section 5 assesses empirically some key predictions of the model using data from Twitter. Section 6 provides additional evidence of attention bartering through Twitter and Instagram user surveys. We discuss some implications of our results in Section 7, and we conclude in Section 8.

2 How Twitter works

Twitter (also known as X) is a platform that allows its users to generate content, called tweets, and to share this content with other users. Tweets are snippets of text, URLs, images, and videos. Twitter serves as the main motivating example throughout the rest of this paper. We henceforth use “Twitter” in lieu of “social media” for concreteness, but our results should be understood to apply to any social medium that shares the features of Twitter that enable attention bartering.

⁵We take users’ desire for attention as a given, but several papers have explored the bases of this desire (Toubia and Stephen, 2013; Del Vicario, Vivaldo, Bessi, Zollo, Scala, Caldarelli and Quattrociocchi, 2016; Pennycook, Epstein, Mosleh, Arechar, Eckles and Rand, 2021).

A user’s tweets are publicly available, unless she has elected to “protect” them by making her account private. Upon logging into Twitter, users see tweets from others in their “timeline”. Twitter offers two different timelines, and users can switch between them on the homepage. The “For You” timeline displays a stream of all public tweets curated by Twitter’s recommendation algorithm.⁶ The “Following” timeline displays a stream of tweets only from the users a user follows. Twitter allows unilateral following: any user is free to follow any other user.⁷ In addition to the tweets, other content may be displayed on a user’s timeline, including suggestions of users to follow and advertisements.

Twitter displays prominently the number of followers and followees each user has. Figure 1a shows a Groucho Marx parody account, which has about 10,800 followers and 734 followees. Additional content the user may opt to provide is also reported, such as the user’s location. When a user views other users’ profile pages, she can see whether she is following them, and whether they are following her. Figure 1b shows an example of how other users’ profiles are viewed when logged in as a user. In this example, the two users follow one another, indicated by the “Following” and “Follows you” tags.

A user can interact with any tweet that is visible to her by “liking” it, “commenting” on it by generating another tweet, or “retweeting” it—which makes the original tweet appear on the timeline of her followers. The user who generated the original tweet receives notifications about the interaction other users have with her tweets. Figure 1c shows a tweet that has received 14 comments, 46 retweets, and 198 likes.

3 A model of social media production and consumption

In this section, we present a model of social media link formation. Each pair from a continuum of (possibly heterogeneous) users decides whether to form bilateral links with each other, which we interpret as attention bartering. After all bilateral bartering decisions are made, users decide whom to follow—a unilateral link—as well. Thus, the coexistence of consumption and attention utilities gives rise to two distinct categories of followers: organic followers who follow a user unilaterally, and reciprocal followers who barter with her.

Individual incentives give rise to a global network structure among the users. In deciding whether to follow unilaterally someone with whom she does not barter, a user considers only the direct consumption value of seeing that user’s content. In contrast, disengaging from a bartering relationship leads to a change in consumption utility, a decrease in attention utility, and a decrease in monitoring or engagement costs. Our paper examines this trade-off for a

⁶Twitter uses a variety of signals such as topic relevance, how popular a tweet is, whether the user is following the author of the tweet, and how the user’s followees are interacting with it to rank and recommend tweets. For more details, see <https://help.twitter.com/en/using-twitter/twitter-timeline>.

⁷Social media platforms Instagram, Youtube, and TikTok also allow unilateral following. However, on some social media, such as Facebook and LinkedIn, all relationships are bilateral.

Figure 1: Elements of Twitter’s information provision.

(a) The profile page of a twitter user.



(b) Features of the relationship of two Twitter users.



(c) Features of a tweet.



Notes: These screenshots depict core elements of Twitter’s information provision. The top panel is the home page of a Groucho Marx parody account. The numbers next to “Following” and “Followers” denote the number of followees and followers of this user. The middle panel is the homepage of The Friars Club account. The page indicates prominently that “The Friars Club” follows the user who is viewing this homepage, and that the user follows The Friars Club, creating a bidirectional following relationship. Note that this relationship may have been either formed organically (with each separately deciding to follow the other), or may be a result of the two users’ attention bartering. The bottom panel is a tweet. The number of comments, retweets, and likes associated with that tweet, appear below the content of the tweet.

user, and how it shapes the network of social media relationships.

3.1 The setting

Our setting is given by a space $\Theta \subseteq \mathbb{R}^d$ (where we can understand a user's type as a list of different attributes);⁸ a probability distribution μ on Θ ; functions $c, b : \Theta \rightarrow \mathbb{R}_{++}$; a function $u : \Theta^2 \rightarrow \mathbb{R}$ such that each $\hat{\theta} \in \Theta$ has $u(\theta, \hat{\theta})$ and $u(\hat{\theta}, \theta)$ nonzero for almost every θ ;⁹ and a strictly concave, increasing function $A : [0, 1] \rightarrow \mathbb{R}$ that is differentiable on $(0, 1]$ and continuous at 0.¹⁰ The interpretation is that a unit mass of Twitter users have publicly observable types $\theta \in \Theta$, distributed according to μ . A given user of type θ derives a net marginal consumption benefit of $u(\theta, \hat{\theta})$ from following a user of type $\hat{\theta}$, pays a constant marginal monitoring/engagement cost of $c(\theta)$ from bartering attention with another user, and derives an attention benefit of $b(\theta)A(\tilde{m})$ from having \tilde{m} followers.¹¹ The possibility that u can take negative values is a reduced-form representation of the fact that a user's total attention (shared between social media and other activities) is scarce, and hence crowding out can be costly.

The general model allows for Twitter users to be flexibly heterogeneous in how much they enjoy reading each others' tweets. If $u(\theta, \hat{\theta}) \geq u(\tilde{\theta}, \hat{\theta})$ for every type $\hat{\theta}$, then type θ users enjoy reading others' content more than type $\tilde{\theta}$ users do; if $u(\hat{\theta}, \theta) \geq u(\hat{\theta}, \tilde{\theta})$ for every type $\hat{\theta}$, then type θ users produce content that others enjoy consuming more than type $\tilde{\theta}$ users' content; and if $u(\theta, \hat{\theta})$ is a decreasing function of types θ and $\hat{\theta}$'s distance, then users prefer to read content from like-minded people. That A is increasing means users derive a benefit from having more followers, and its concavity reflects diminishing returns from receiving attention. We allow for users to be heterogeneous in how much they value attention rather than consumption, as reflected by the user-specific quantity $b(\theta)$. Finally, we assume that users also incur a cost $c(\cdot)$ for every bartering relationship they maintain. This cost can be interpreted as a cost of maintaining engagement in a relationship, or as a cost of monitoring whether the partner does in fact continue to follow the user; our model allows for this cost to be heterogeneous.

To build up to our equilibrium concept, we invest in some useful notation.

Notation 1. Let \mathcal{F}_1 denote the set of measurable functions $\Theta^2 \rightarrow [0, 1]$. Let \mathcal{F}_2 denote the

⁸We assume Θ , and any other set we refer to, is Borel; and we use the word “function” as a short-hand for a bounded Borel function. We use $\mathbb{E}_{\hat{\theta}}[\cdot]$ as a short-hand for $\int[\cdot] d\mu(\hat{\theta})$ whenever this expectation exists.

⁹This condition ensures optimal *unilateral* following decisions are essentially unique, and so facilitates our focus on reciprocal following.

¹⁰Note, these assumptions imply A is continuous on $[0, 1]$ and continuously differentiable on $(0, 1]$, with $\lim_{\tilde{m} \searrow 0} A'(\tilde{m}) = \lim_{\tilde{m} \searrow 0} \frac{A(\tilde{m}) - A(0)}{\tilde{m}} =: A'(0) \in [0, \infty]$.

¹¹At the expense of additional notation, one can easily extend the model to allow a user's attention benefit to depend on the types of users who follow her, rather than simply the total quantity of followers. One could also easily extend the model to allow the attention benefit to depend on endogenous features, such as the number of followers a follower has.

set of $f_2 \in \mathcal{F}_1$ such that $f_2(\theta, \hat{\theta}) = f_2(\hat{\theta}, \theta)$ for every $\theta, \hat{\theta} \in \Theta$. Given any $(f_1, f_2) \in \mathcal{F}_1 \times \mathcal{F}_2$ and $\theta \in \Theta$, let

$$\begin{aligned} m(\theta|f_1, f_2) &:= E_{\hat{\theta}} \left\{ f_2(\theta, \hat{\theta}) + [1 - f_2(\theta, \hat{\theta})] f_1(\hat{\theta}, \theta) \right\}, \\ V(\theta|f_1, f_2) &:= E_{\hat{\theta}} \left\{ f_2(\theta, \hat{\theta}) \left[u(\theta, \hat{\theta}) - c(\theta) \right] + [1 - f_2(\theta, \hat{\theta})] f_1(\theta, \hat{\theta}) u(\theta, \hat{\theta}) \right\} \\ &\quad + b(\theta) A(m(\theta|f_1, f_2)). \end{aligned}$$

Let us interpret the above objects. First, $f_1 \in \mathcal{F}_1$ specifies unilateral following behavior: $f_1(\theta, \hat{\theta})$ is the probability that a user of type θ unilaterally follows one of type $\hat{\theta}$, *conditional on* them not engaging in attention bartering. Next, $f_2 \in \mathcal{F}_2$ specifies reciprocal following behavior: $f_2(\theta, \hat{\theta})$ is the probability with which a user of type θ and one of type $\hat{\theta}$ attention barter—which therefore is assumed symmetric. Given a pair (f_1, f_2) of such functions, $m(\theta|f_1, f_2)$ records the number of followers a user of type θ has, whereas $V(\theta|f_1, f_2)$ is such a user's payoff.

Remark 1. Notice, f_1 summarizes hypothetical unilateral following decisions conditional on no bartering. So types θ and $\hat{\theta}$ attention barter if $f_2(\theta, \hat{\theta}) = 1$; and type θ unilaterally follows $\hat{\theta}$ if $f_2(\theta, \hat{\theta}) = 0$ and $f_1(\theta, \hat{\theta}) = 1$. In particular, if $f_2(\theta, \hat{\theta}) = 0$ and $f_1(\theta, \hat{\theta}) = f_1(\hat{\theta}, \theta) = 1$, then both types follow each other but (as nobody pays the cost c) we do not interpret their relationship as one of attention bartering.

This specification reflects a specific timing interpretation: users make unilateral following decisions taking bartering choices as given, and they make bartering decisions anticipating that unilateral following decisions will be made optimally.

We interpret any $f_1 \in \mathcal{F}_1$ as a directed social network and $f_2 \in \mathcal{F}_2$ as an undirected social network with the same set of users as nodes. Our interest is in which pairs of such networks can arise in an equilibrium, which we formalize below.

Definition 1. An *equilibrium* is a pair $(f_1, f_2) \in \mathcal{F}_1 \times \mathcal{F}_2$ such that

1. Every $\theta, \hat{\theta} \in \Theta$ have

$$f_1(\theta, \hat{\theta}) = \begin{cases} 1 & : u(\theta, \hat{\theta}) > 0 \\ 0 & : u(\theta, \hat{\theta}) < 0. \end{cases}$$

2. Every $\theta \in \Theta$ has $V(\theta|f_1, f_2) \geq V(\theta|f_1, \tilde{f}_2)$ for each $\tilde{f}_2 \in \mathcal{F}_2$ such that

- Every $\hat{\theta}, \tilde{\theta} \in \Theta \setminus \{\theta\}$ have $\tilde{f}_2(\hat{\theta}, \tilde{\theta}) = f_2(\hat{\theta}, \tilde{\theta})$;
- Every $\hat{\theta} \in \Theta$ with $\tilde{f}_2(\theta, \hat{\theta}) > f_2(\theta, \hat{\theta})$ has

$$[1 - f_1(\hat{\theta}, \theta)]u(\hat{\theta}, \theta) - c(\hat{\theta}) + [1 - f_1(\hat{\theta}, \theta)]b(\hat{\theta})A'(m(\hat{\theta}|f_1, f_2)) > 0.$$

The interpretation of the equilibrium is that users make their reciprocal following decisions taking into account anticipated unilateral following decisions. The first condition says that a user unilaterally follows another (conditional on not having an attention bartering relationship) if the flow consumption benefit of following them is positive, and not if it is negative. Because a user’s unilateral following decisions do not affect her own follower count, and the rest of her objective is additively separable across other users, we have written this constraint in separable form. The second condition captures the optimality of a user’s attention bartering behavior. When considering possible deviations for the user, the definition assumes that (i) users cannot change relationships that do not affect themselves, (ii) users can unilaterally break ties to end any bartering relationship, and (iii) users can only instigate a new bartering relationship with partners who acquiesce.¹² That bartering relationships can be formed with bilateral consent and broken unilaterally is in the spirit of *pairwise stability* (Jackson and Wolinsky, 1996).

Table 1: Model notation

Term	Definition
θ	User type
$u(\theta, \hat{\theta})$	Marginal consumption benefit for user type θ derived from following type $\hat{\theta}$
$c(\theta)$	Monitoring cost of attention bartering for user type θ
$b(\theta)A(\tilde{m})$	Attention benefit for user type θ from having \tilde{m} followers, where A is a strictly concave, increasing function; $b = 1$ in the vertical model
$f_1(\theta, \hat{\theta})$	Probability that user type θ unilaterally follows user type $\hat{\theta}$, conditional on no attention bartering
$f_2(\theta, \hat{\theta})$	Probability that user type θ and user type $\hat{\theta}$ attention barter; f_2 is symmetric—i.e., $f_2(\theta, \hat{\theta}) = f_2(\hat{\theta}, \theta)$
$\kappa(\theta, \hat{\theta})$	User type θ ’s normalized marginal cost of attention bartering with user type $\hat{\theta}$
$A'(m(\theta))$	User type θ ’s normalized marginal attention benefit of having an additional follower

3.2 Characterizing equilibrium

The following statistic of users types will turn out to be useful in our analysis.

¹²Consistent with the interpretation of individual users being “small,” we express the would-be partner’s willingness to form a new link through a first-order condition: new attention bartering entails a marginal cost, a marginal consumption utility, and a marginal follower increment, with the latter two only applying net of how this dyad would behave if they did not barter.

Notation 2. Let $\kappa : \Theta^2 \rightarrow \mathbb{R}_{++}$ be the function given by

$$\kappa(\theta, \hat{\theta}) := \frac{c(\theta) + \max\{0, -u(\theta, \hat{\theta})\}}{b(\theta)}.$$

The function $\kappa(\theta, \cdot)$ constitutes a lower-dimensional description of type θ 's preferences; combining it with θ 's follower count will allow us to describe the marginal rate of substitution between following a given type and being followed by it, which is relevant to bartering decisions.

The following theorem explicitly characterizes equilibria. In many cases of interest (e.g., if $\kappa(\theta, \cdot)$ is atomlessly distributed for every $\theta \in \Theta$), the proposition gives an essentially unique candidate equilibrium compatible with any given candidate follower count $m : \Theta \rightarrow [0, 1]$.

Theorem 1. A pair $(f_1, f_2) \in \mathcal{F}_1 \times \mathcal{F}_2$ is an equilibrium if and only if:

1. Every $\theta, \hat{\theta} \in \Theta$ have:

$$f_1(\theta, \hat{\theta}) = \begin{cases} 1 & : u(\theta, \hat{\theta}) > 0 \\ 0 & : u(\theta, \hat{\theta}) < 0; \end{cases}$$

2. Defining $m(\cdot) := m(\cdot | f_1, f_2)$, every $\theta \in \Theta$ and almost every $\hat{\theta} \in \Theta$ have

$$f_2(\theta, \hat{\theta}) = \begin{cases} 1 & : u(\theta, \hat{\theta}) \leq 0, u(\hat{\theta}, \theta) \leq 0, \\ & \kappa(\theta, \hat{\theta}) < A'(m(\theta)), \text{ and } \kappa(\hat{\theta}, \theta) < A'(m(\hat{\theta})) \\ 0 & : u(\hat{\theta}, \theta) > 0 \text{ or } \kappa(\theta, \hat{\theta}) > A'(m(\theta)). \end{cases}$$

The proof of the theorem, and all other proofs, can be found in Appendix A.

The above result clarifies which features of individual preferences determine who follows whom. Intuitively, $\kappa(\theta, \hat{\theta})$ captures type θ 's normalized marginal cost of bartering with user $\hat{\theta}$, which is compared to the normalized attention benefit $A'(m(\theta))$ of having an additional follower.

3.3 Consequences of equilibrium attention bartering

Let us highlight an interpretable property of all equilibria. The following corollary is an immediate consequence of Theorem 1.

Corollary 1. In any equilibrium (f_1, f_2) , each $\theta \in \Theta$ and almost every $\hat{\theta} \in \Theta$ with $u(\hat{\theta}, \theta) > 0$ have $f_2(\theta, \hat{\theta}) = 0$.

The corollary results in an immediate welfare ranking between the attention bartering economy and a counterfactual arrangement in which bartering is impossible. Indeed, the corollary

says that (up to measure-zero exceptions) any user who *would* unilaterally follow type θ —a decision which would remain the same even if attention bartering were not available—does not barter with type θ users in equilibrium. Said differently, attention bartering never crowds out unilateral following in our model.

Another byproduct of Theorem 1 is a criterion to rank users unambiguously. Although equilibrium may not be unique (as the next subsection shows) because users' incentives cannot be determined without considering the entire network, the following proposition shows we can still compare the equilibrium outcomes of some users.

Proposition 1. *If $\theta, \tilde{\theta} \in \Theta$ have $u(\hat{\theta}, \theta) > u(\hat{\theta}, \tilde{\theta})$ and $\kappa(\theta, \hat{\theta}) \leq \kappa(\tilde{\theta}, \hat{\theta})$ for almost every $\hat{\theta} \in \Theta$, then every equilibrium (f_1, f_2) has $m(\theta|f_1, f_2) \geq m(\tilde{\theta}|f_1, f_2)$.*

Intuitively, if a user of type θ generates greater consumption benefit than one of type $\tilde{\theta}$, then more users will be willing to follow her unilaterally; and among those who do not, more will be willing to attention barter with her. Hence, the former type will have more choices available than the latter. Furthermore, if users of type θ have a lower marginal cost of bartering with anyone, then they cannot have fewer followers than users of type $\tilde{\theta}$ in equilibrium—because bartering with some of $\tilde{\theta}$'s willing barterers is a profitable deviation.

3.4 Homogeneous users

In this subsection, we consider the case in which users are homogeneous, i.e., the functions u, c, b are all constant functions (with $u \neq 0$ by hypothesis), and hence κ is constant too. This special case is useful for plainly demonstrating some of the forces in our model. The following proposition is an immediate consequence of Theorem 1.

Corollary 2. *With homogeneous users, equilibrium is essentially characterized as follows:¹³*

- (i) *If $u > 0$, everybody follows everybody unilaterally.*
- (ii) *If $u < 0$ and $\kappa \leq A'(1)$, everybody attention barterers with everybody.*
- (iii) *If $u < 0$ and $\kappa \geq A'(0)$, nobody follows anybody.*
- (iv) *If $u < 0$ and $\kappa \in (A'(1), A'(0))$, then let $m^* := (A')^{-1}(\kappa)$. In this case, nobody follows anybody unilaterally, everybody attention barterers with at most m^* others, and everybody with strictly fewer than m^* followers barterers with each other.*

The second case highlights an efficiency gain enabled by attention bartering. When $\kappa \leq A'(1)$, the utilitarian surplus always increases with additional following, even if the cost c must be

¹³That is, any equilibrium can be modified on a measure-zero set to some $(f_1, f_2) \in \mathcal{F}_1 \times \mathcal{F}_2$ that has these properties, and any $(f_1, f_2) \in \mathcal{F}_1 \times \mathcal{F}_2$ that has these properties can be modified on a measure-zero set to yield an equilibrium.

borne by the follower. However, if this condition holds while $u < 0$ (which happens for some parameter values), then users cannot commit to follow each other without attention bartering, because they do not internalize the positive externality they exert on their followees. In this case, no following would happen if attention bartering were not feasible. However, with attention bartering, everybody is followed by everybody, partially restoring efficiency.

The fourth case highlights the global nature of the equilibrium conditions, even while following decisions are local to a specific pair of users. Specifically, it highlights that whether or not two users face the “double coincidence of wants” needed for attention bartering depends on their overall “popularity,” which is a global feature of the equilibrium. The following parametric example demonstrates how this case can render equilibrium outcomes non-unique.

Consider the parametric example with $\Theta = [0, 1]$, μ uniform, constant functions $u = -1$ and $c = b = 1$, and $A(m) = 3.2\sqrt{m}$. In this example, $u < 0$, $A'(m) = 1.6/\sqrt{m}$, and $\kappa = 2$. Hence, a user’s targeted number of followers (as given in the fourth case of Corollary 2) is $m^* = 0.64 \in (0.5, 1)$. Observe some equilibrium exists in which nobody follows anybody unilaterally; types below 0.64 barter with each other, and types above 0.64 barter with each other. Intuitively, the types on the left are doing exactly the amount of attention bartering they would like; and the types on the right would like to be doing more, but are already following everybody who would willingly reciprocate. In this equilibrium, the types on the left have more followers, and a higher payoff than those on the right. Moreover, by homogeneity, an alternative equilibrium exists in which $\Theta_1 = (0.36, 1]$ forms a clique and $\Theta \setminus \Theta_1$ forms a clique. In particular, neither equilibrium payoffs nor equilibrium follower counts are unique in general, and there can exist equilibria that are not Pareto ranked.

One can find other equilibria that Pareto dominate the above. For instance, we could take $f_1 = 0$ and $f_2(\theta, \hat{\theta}) = \mathbf{1}_{\rho(\theta, \hat{\theta}) \leq 0.32}$, where the metric ρ on Θ is given by $\rho(\theta, \hat{\theta}) := \min\{|\theta - \hat{\theta}|, 1 - |\theta - \hat{\theta}|\}$.¹⁴ In this case, every type attention barterers with types within a small enough ρ -distance, and the equilibrium entails efficient (indeed, Pareto-dominant) bartering: it maximizes every user’s payoff subject to the constraint that no unilateral following happens when $u < 0$.

4 Vertically differentiated users

In this section, we consider a special case of our model, which we term the **vertical model**. The substantive special feature of the vertical model is that users are purely vertically differentiated.¹⁵ As we will show, this special case of the model has an essentially unique

¹⁴If we identify $[0, 1]$ with a circle by mapping θ to $e^{2\pi i \theta}$, then ρ captures how far apart two types are, as a fraction of the circumference.

¹⁵For simplicity, we cast this model with users’ sole heterogeneous attribute being the consumption benefit they yield to those who follow them, but as the analysis should make clear (in light of Theorem 1), this particular functional form is not central to the results, and could be replaced with an appropriate single-

equilibrium network and generates a rich set of testable predictions that are not wedded to a particular parameterization. Accordingly, this version of the model is especially well suited to empirical analysis; we therefore use it as guide in Section 5.

Informally, users' types are linearly ordered in the vertical model. We explicitly solve for the resulting equilibrium of the attention bartering economy and show it is essentially unique. The equilibrium takes a core-periphery structure, with the periphery partitioned into a number of vertically differentiated "clubs" that are isolated from one another.

4.1 The setting for the vertical model

Assume $c > 0$ is a constant, $b = 1$ is a constant, $\Theta \subset \mathbb{R} \setminus \{0\}$ is bounded, the population distribution μ has a continuous cumulative distribution function $G : \mathbb{R} \rightarrow [0, 1]$, and $u(\theta, \hat{\theta}) = \hat{\theta}$ for every $\theta, \hat{\theta} \in \Theta$. In this stylized special case of our model, we can interpret $\theta \in \Theta$ as a user's **ability**.

To describe the equilibrium network, the notation defined in the following claim is useful.

Claim 1. *Suppose $A'(0) \geq c$, and let $\theta_0 := 0$. Recursively, for $n \in \mathbb{N}$, a unique $\theta_n \in [-\infty, \theta_{n-1}]$ exists with $A'(G(\theta_{n-1}) - G(\theta_n)) = c - \theta_n$. Moreover, $\theta_\infty := \lim_{n \rightarrow \infty} \theta_n = c - A'(0) \in [-\infty, 0]$.¹⁶*

4.2 Characterizing equilibrium for the vertical model

We now present our main result for the vertical model. The following proposition explicitly describes an equilibrium, and shows it is essentially unique.

Proposition 2. *Consider the vertical model, and let $\{\theta_n\}_{n=0}^\infty$ be as defined in Claim 1 in the case that $A'(0) \geq c$. The pair $(f_1^*, f_2^*) \in \mathcal{F}_1 \times \mathcal{F}_2$ with each $\theta, \hat{\theta} \in \Theta$ having*

$$\begin{aligned} f_1^*(\theta, \hat{\theta}) &= \mathbf{1}_{\hat{\theta} > 0} \\ f_2^*(\theta, \hat{\theta}) &= \begin{cases} \mathbf{1}_{(\theta, \hat{\theta}) \in \bigcup_{n=1}^\infty (\theta_n, \theta_{n-1}]^2} & : A'(0) \geq c \\ 0 & : A'(0) \leq c \end{cases} \end{aligned}$$

is an essentially unique equilibrium.

If $A'(0) \leq c$, then bartering with another user is never worthwhile on the margin, unless following unilaterally the same user is desirable. No attention bartering happens in this case, but some unilateral following can still happen: users of type $\hat{\theta} > 0$ are **stars** who everybody follows unilaterally and who barter with nobody.

crossing condition.

¹⁶Note that the cutoffs $\{\theta_n\}_{n \in \mathbb{Z}_+ \cup \{\infty\}}$ might not belong to Θ . Nevertheless, they are helpful in describing the equilibrium.

The more interesting case is that with $A'(0) > c$. In this case, users of type $\hat{\theta} > 0$ are still stars who everybody follows unilaterally—including the other stars—and who attention barter with nobody. In addition, some attention bartering happens, and it takes a “club” form. Specifically, each of the disjoint sets $\{(\theta_n, \theta_{n-1}]\}_{n=1}^{\infty}$ forms a clique, wherein each pair of users with types in $(\theta_n, \theta_{n-1}]$ barter attention. No bartering happens across clubs, although all club members still follow unilaterally the stars. Finally, any users with type below $\theta_{\infty} = c - A'(0)$ are **lurkers**, who have no followers of either kind. Intuitively, such a user is subject to a Groucho Marx condition: she does not want to be a member of any club that would have her as a member.

The next corollary follows immediately from Proposition 2. It characterizes, in terms of primitives, when the equilibrium has a positive number of stars and when it has a positive number of lurkers.

Corollary 3. *In the vertical model, in equilibrium:*

- (i) *A positive measure of stars exist if and only if $G(0) < 1$.*
- (ii) *A positive measure of lurkers exist if and only if $G(c - A'(0)) > 0$.*

Our characterization also generates clear patterns on how a user’s social connections vary with her ability. Intuitively (and consistent with Proposition 1), higher-ability users have more followers. However, the relationship between ability and number of followees has a more interesting pattern: a followee dip. The following proposition summarizes this and the related testable predictions of the vertical model.

Proposition 3. *In the vertical model, if $G(0) < 1$, any equilibrium essentially has:*

- (i) *The number of followers and follower-to-followee ratio are weakly increasing in user ability. Only stars have follower-to-followee ratios higher than one.*
- (ii) *The number of followees is quasiconcave (that is, single-peaked) in user ability: it weakly increases initially, and then decreases for stars, who follow the same number of users as lurkers.*

In particular, a user’s number of followees is non-monotone in her ability if $A'(G(0)) + \underline{\theta} < c < A'(0)$, where $\underline{\theta}$ is the minimum of the type distribution’s support.
- (iii) *Lurkers and stars do not engage in attention bartering. Every other user’s number of reciprocal followers is positive and weakly increasing in her ability.*
- (iv) *Users comprising an attention bartering club have the same number of followers, followees, and ratios. The size of the clubs—and hence these statistics—increases weakly in user ability among those users who are neither lurkers nor stars.*

4.3 An example

We conclude this section with a parametric example of the vertical model. Suppose types are uniformly distributed on $[\underline{\theta}, \bar{\theta}] \setminus \{0\}$, where $\bar{\theta} \in (0, 1)$ and $\underline{\theta} = \bar{\theta} - 1 \in (-1, 0)$, and A takes the form $A(\tilde{m}) = 2\sqrt{s\tilde{m}}$ for some scale factor $s > 0$. In this case, the recursive definition of $\{\theta_n\}_{n=1}^\infty$ in Claim 1 can be written more explicitly. Recursively, for $n \in \mathbb{N}$, we obtain θ_n as the unique solution to

$$(\theta_{n-1} - \theta_n)(c - \theta_n)^2 = s$$

whenever that solution is above $\underline{\theta}$. In this example, the equilibrium network has some stars, but no lurkers (by Corollary 3). Moreover, by Proposition 3(ii), a user’s number of followers is non-monotone in her ability if $s < (1 - \bar{\theta})(c + 1 - \bar{\theta})^2$.

To illustrate further the equilibrium of our model, we consider in what follows the case of attention benefit $A(m) = 0.5\sqrt{m}$, abilities θ uniformly distributed on $[-0.8, 0.2]$, and cost $c = 0.2$. Figure 2a plots the equilibrium of the counterfactual setting in which attention bartering is impossible. Users cannot barter and hence they follow unilaterally only the stars, obtaining consumption utility $V_0 = \int_0^{0.2} \theta \, d\theta$. Users with ability $\theta \leq 0$ lurk: they are not followed by anyone, they do not generate tweets, and they obtain total utility equal to their consumption utility. In contrast, star users are followed by every user, tweet, and obtain total utility equal to $V_0 + A(1)$.

Figure 2b illustrates the equilibrium of our model (in which attention bartering is possible). Star users are unaffected: they are followed unilaterally by every other user, and they do not engage in reciprocal following. The set of lurkers vanishes, as non-star users can improve their total utilities by forming attention bartering clubs. Seven attention bartering clubs form in this parametric example; the largest comprises users of types $(\theta_1, \theta_0]$ where $\theta_0 = 0$, and the smallest comprises users of types $[\underline{\theta}, \theta_6]$. Note that the extent of reciprocal following—the size of these clubs—becomes smaller for clubs consisting of lower-ability users, as bartering with those users is costlier (and so it requires the higher marginal value of attention that a smaller follower count affords). Club members are better off in terms of total utility but their consumption utility decreases, as they follow users who generate a net negative consumption value. Higher-ability users belong to more populous clubs, and hence incur higher monitoring costs in addition to the consumption cost of attention bartering. We also depict these predictions in Figure 4 and Figure 5, where we compare the model predictions with estimates from #EconTwitter data—the subject of the paper’s next section.

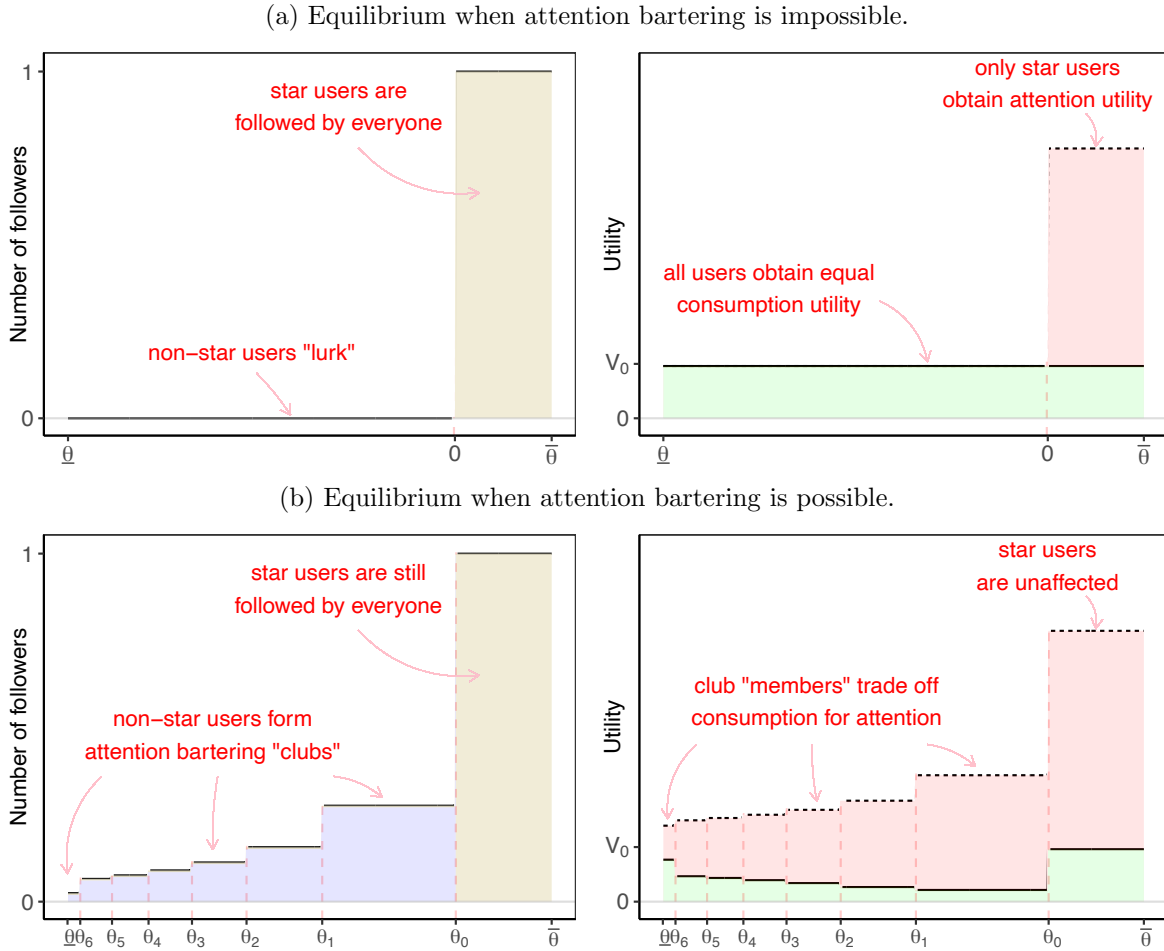
5 Assessing model predictions with #EconTwitter data

We next assess empirically some core assumptions and predictions of our model. To narrow the focus of this exercise, we restrict attention to the vertical specialization of our model in

which users have homogeneous tastes but are heterogeneous in the value others derive from their content.¹⁷ To that end, we collected data from #EconTwitter, a Twitter community comprising professional and amateur users who tweet mostly about economics, and often fol-

¹⁷Throughout this section, as we focus wholly on the vertical model of Section 4, we refer to it (for brevity) simply as our model. Recall, any specification of the vertical model generates an essentially unique equilibrium. Although the qualitative predictions of this equilibrium we test are not specific to the parameterization, all figures representing the model adopt the same parameterization as in Figure 2.

Figure 2: Following, consumption, and attention in the vertical model.



Notes: This figure plots the equilibria with and without attention bartering, for the case of uniformly distributed types on the interval $[\underline{\theta}, \bar{\theta}] = [-0.8, 0.2]$, $A(m) = 0.5\sqrt{m}$, and $c = 0.2$. In all panels, the x-axis depicts user types, the yellow-shaded area depicts the number of organic followers, the blue-shaded area depicts the number of reciprocal followers, the green-shaded area depicts the consumption utility minus monitoring costs, and the red-shaded area depicts the attention utility. Panel 2a plots the equilibrium when attention bartering is impossible. Without bartering, only stars obtain attention utility. Panel 2b plots the the equilibrium when attention bartering is possible. Stars are unaffected, but non-star users form attention bartering clubs emerge, and they trade off decreases in consumption utility for increases in attention utility. Figure 4 and Figure 5 depict additional quantities for the attention bartering equilibrium.

low each other.¹⁸ Our rationale for focusing on #EconTwitter is to obtain data that matches closely the vertical setting, that is, data on users with far less horizontal differentiation than one would expect from a random sample of Twitter users.

Let us emphasize the nature of our empirical exercise. We are not fitting the #EconTwitter data to our model and estimating parameters of the latter. Rather, we are comparing the qualitative patterns in the #EconTwitter data—specifically, in how network statistics vary with a proxy for user ability—to the patterns predicted by our model. For example, our model predicts a “followee dip” at the highest levels of user ability, and we find such a dip in the #EconTwitter data.

Our exercise is somewhat analogous to comparing the signs of coefficients in a regression to the comparative statics of a model. However, our approach is more restrictive: we make such comparisons at many ability levels, and the signs of some model predictions change with the ability levels. Returning to the followee dip example, the number of followees a user has is non-monotone in her number of followers: it increases with ability up until the level at which a user attracts organic followers, at which point the sign changes and the followee count drops.

5.1 Data collection and sample definition

Because Twitter has no direct “tag” indicating that a user is an economist, we adopted a heuristic process based on Twitter’s list feature. Lists are groups of Twitter users, created and curated by individual users in order to organize tweets thematically. Lists can be made private or public, and can be shared with other users. Importantly, adding a user to a list does not affect follower or followee counts, and hence prospective followers cannot see the number of lists to which a user has been added.¹⁹

We began by collecting data on users comprising RePEc’s “Economists on Twitter” list. We then collected data on users followed by at least 3 users on the RePEc list, for a total of 68,592 users. The data features for each user include the unique identifiers of her followers and followees, the number of tweets she has produced, the number of tweets she has liked, her tenure on the platform, and the number of lists by other users that on which she appears. All data were collected using Twitter’s public API.²⁰

¹⁸On Twitter, users may prepend the hashtag character to relevant keywords or phrases within their tweets, as a form of tagging that enables cross-referencing of content sharing a theme or subject. Twitter communities often derive their names from recurring hashtags that their members use to broadcast messages to the community. For more details, see <https://help.twitter.com/en/using-twitter/how-to-use-hashtags>, and <https://en.wikipedia.org/wiki/Hashtag>.

¹⁹For more details on Twitter lists, see <https://help.twitter.com/en/using-twitter/twitter-lists>.

²⁰The RePEc list can be found at <https://ideas.repec.org/i/etwitter.html>. The official documentation describing the features of data obtained through Twitter’s API can be found at <https://developer.twitter.com/en/docs/twitter-api/v1/data-dictionary/object-model/user>.

We refined the data sample as follows. First, we removed “inactive” users, defined as users with fewer than 25 followers, fewer than 10 followees, or fewer than 50 tweets. Second, we removed “superstar” users, defined as users with more than 200,000 followers. This restriction removes few economist outliers, but many non-economists who are particularly popular with #EconTwitter users.²¹ Third, as a crude measure to eliminate some “spam” accounts from our data, we removed users with more than 5,000 followees. This restriction drops only a small number of users, and is consistent with the Twitter-imposed basic followee limit.²² The final sample consists of 55,966 #EconTwitter users, whose followers and followees consist of 26,284,663 unique users.

5.2 Followers, followees, and ratios

Figure 3 reports the distributions of followers, followees, and follower-to-followee ratios (henceforth, “ratios” for brevity) of #EconTwitter users. We plot a Gaussian kernel density estimate of the distribution of each outcome, with the bandwidth selected using Silverman’s rule of thumb (Silverman, 1986). The red vertical lines depict the median (solid) and the mean (dashed) of each distribution.

On average, an #EconTwitter user has 10,304 followers, 1,273 followees, and a ratio of 13.98, whereas the median user has 2,004 followers, 913 followees, and a ratio of 2.05.²³ Follower counts are highly right-skewed, with some star users that attract vast numbers of followers. This feature is consistent with producers facing no distribution costs and being vertically differentiated, with the “best” users having large audiences. In contrast, the distribution of users’ followees is substantially less right-skewed, and does not have a long tail. This feature is consistent with consumers having scarce attention.


Despite this follower/followee distinction—and some star users having enormous ratios—follower-to-followee ratios are clustered around one, consistent with Proposition 3(i). This clustering of ratios near one has been documented in previous work spanning several social media, and a relatively long period of time (Kwak et al., 2010; Myers et al., 2014; Sadri et al., 2018). Interestingly, our “clubs”—where all users follow each other—have ratios slightly less than one, with the denominator inflated by the “stars” those users follow. If the network has few stars, our model predicts large numbers of users with ratios of approximately one.

²¹For example, Barack Obama, Shakira, and Britney Spears are the three most followed non-economists who are followed by 3 or more users in the RePEc list. Among economists in the RePEc list, the “superstar” restriction resulted in Paul Krugman, Xavier Sala-i-Martin, Joseph Stiglitz, and Alejandro Gaviria being dropped from the sample. Erik Brynjolfsson (barely) made it into our refined data set.

²²For more details, see <https://help.twitter.com/en/using-twitter/twitter-follow-limit>.

²³In Appendix B, we report estimates computed on the entire sample. Compared to the average Twitter user, #EconTwitter users have on average more followers and followees, and higher ratios: the median user in the entire data set has 266 followers, 524 followees, and a ratio equal to 0.52.

Figure 3: Distributions of followers, followees, and ratios of #EconTwitter users



Notes: This figure reports kernel density estimates of the probability density function of the distributions of followers, followees, and follower-to-followee ratios for #EconTwitter users. For each facet, the vertical red lines depict the median (solid) and the mean (dashed) of the corresponding distribution, and the kernel bandwidth is selected using Silverman’s rule of thumb (Silverman, 1986). See Appendix B for density estimates of the same statistics for all users.

5.3 Constructing a proxy for user ability using lists

A key distinguishing feature of our model is its joint predictions concerning ability and outcomes. However, the “abilities” of Twitter users cannot be observed directly. As a proxy for user ability, we use the number of lists to which a user has been added by other users.

Twitter lists have several properties that are useful for constructing a proxy for user ability. First, adding a user to a list indicates interest in her tweets, without increasing the enlister’s followee count. Second, the number of lists to which a user has been added is not observable by her would-be followers, and hence list-adding is not readily bartered in the same manner as attention.

One limitation of the list count is that users with greater tenure on the platform are likely to have been added to more lists. To remove some of the tenure dependency effect, we residualize the log-number of lists to which a user has been added, by partialing out the log-number of tweets that the same user has generated.²⁴ We compute the quantile each user occupies in our data with respect to this residualized number of lists, and refer to this number as “user ability” throughout the rest of this section.

²⁴We do not control for tenure using the time since a user joined the platform, because several users initially tweet rarely, or they start tweeting long after they join the platform.

5.4 User ability and network statistics

Our model makes sharp predictions about users’ network statistics in the attention bartering equilibrium. In Figure 4, we compare these predictions to estimates obtained from the #EconTwitter data. Each row depicts a separate network statistic. The left-column facets depict the empirical estimates as a function of the list-defined user ability. Each point represents the mean value for users belonging to the same ability percentile, and a 95% confidence interval is plotted around each point estimate. The right-column facets depict the equilibrium predictions as a function of user ability, using the model parameterization of Figure 2. The model parameters were chosen for ease of exposition—what matters for our empirical exercise is the shape of the equilibrium predictions, which are not specific to the parameterization.

Facet (a) plots the number of users’ followers—the number of users following each user. Followers increase in users’ abilities, suggesting that our choice of proxy is reasonable. Consistent with our modeling assumption and with Proposition 3(i), high-ability users attract a far larger number of followers than low-ability users.

Facet (b) plots the number of followees—the number of users followed by each user. Followees increase initially in user ability, but at a much slower rate than the number of followers, and then decrease at high ability levels.²⁵ This “followee dip” is consistent with Proposition 3(ii), and supports our model’s feature that following users is costly in terms of consumption, rather than just being limited by awareness or opportunity.

Follower-to-followee ratios are reported in facet (c). These ratios are increasing in user ability, but at a slower rate than the number of followers: users within broad ability ranges have similar ratios, consistent with Proposition 3’s parts (i) and (iv).²⁶

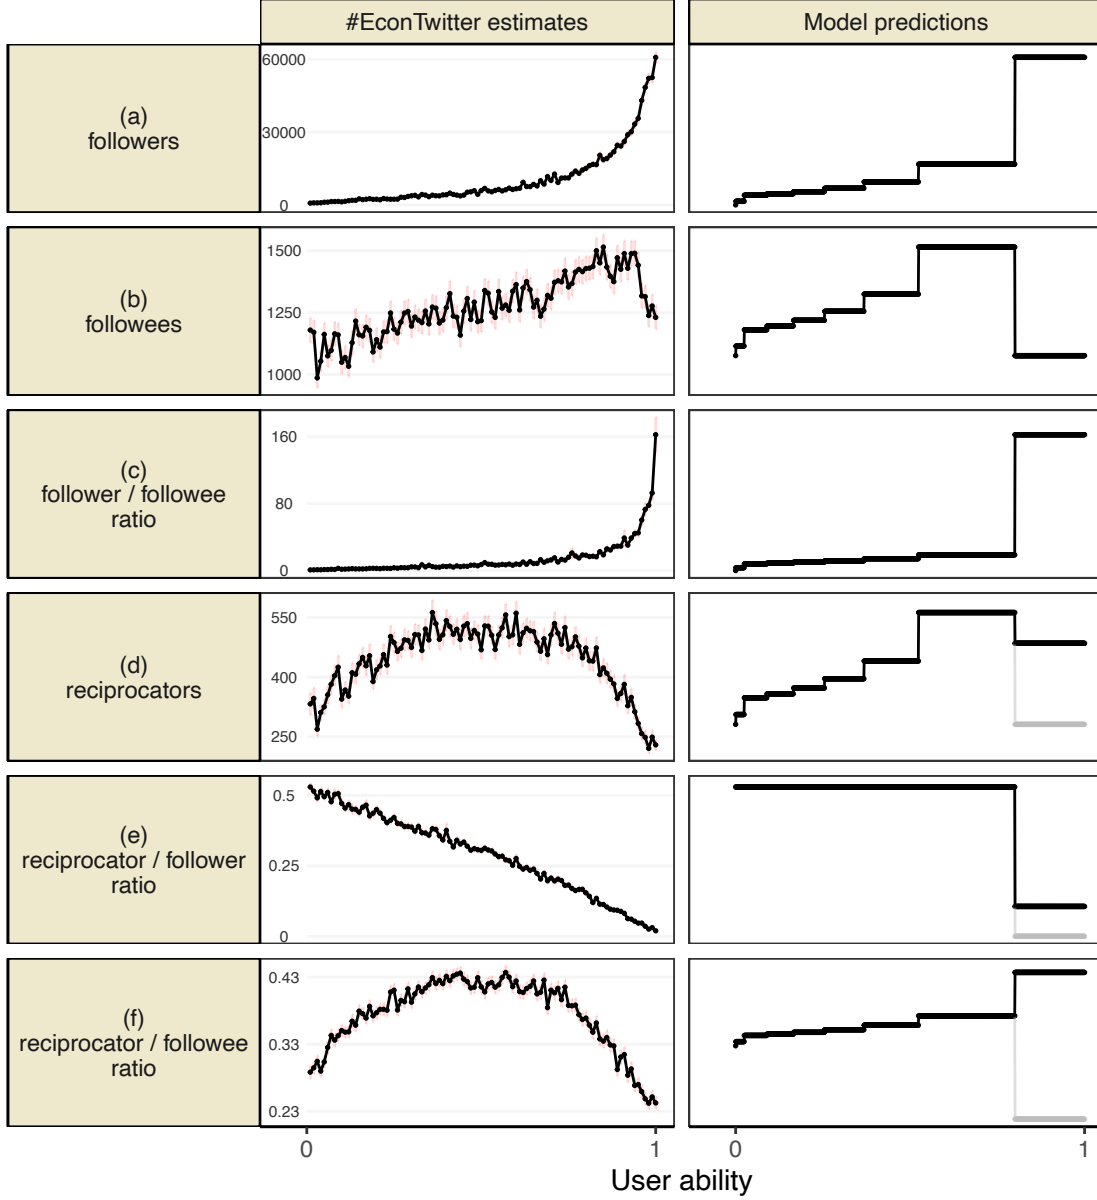
5.5 Reciprocation and attention bartering

We cannot observe whether a relationship is truly reciprocal in the sense of our model—we do not know what would happen if one party were to unfollow the other. As such, we cannot verify from a network snapshot whether a bilateral relationship results from the two users following each other organically, or from attention bartering. Said differently, all we directly observe in the data is “gross reciprocation” as captured by any bidirectional following relationship. On the right-hand-side of facets (d), (e), and (f) of Figure 4, black lines depict such gross reciprocation, whereas gray lines depict true reciprocal following in the form of attention bartering.

²⁵For example, economists Susan Athey and David Autor belong to the highest ability tier in our data, but only follow a below-median 875 and 176 users, respectively.

²⁶At high ability levels, ratios seemingly become more dispersed within the same ability level. One possible explanation is that high-ability users, as captured by our proxy, comprise both (1) star users with many followers and few followees, (2) users with large numbers of both followers and followees. The latter may be users who have attained visibility due to attention bartering and have been subsequently added to many lists—illustrating a limitation of our ability proxy.

Figure 4: #EconTwitter estimates and model predictions of network statistics.



Notes: This figure reports empirical estimates and theoretical predictions for networks statistics as a function of normalized user ability. On the left column, we plot estimates obtained from the #EconTwitter data. Each point represents the mean value (y-axis) for users belonging to the same ability percentile (x-axis), and a 95% confidence interval is shown around each point estimate. See Section 5.3 for details on the definition of #EconTwitter users' abilities. On the right column, we plot model predictions for the same quantities. The illustrated attention bartering equilibrium predictions are for the case of uniformly distributed types on $[-0.8, 0.2]$, $A(m) = 0.5\sqrt{m}$, and $c = 0.2$; abilities are represented by their quantiles, and all dependent variables are scaled to be displayed alongside the empirical estimates. For panels (d) - (f), the black line depicts network statistics assuming that two users following each other organically are reciprocators, and the gray line depicts network statistics assuming that two users reciprocate only when they attention barter with one another. See Figure 2b for an alternative representation of the same attention bartering equilibrium.

Facet (d) in Figure 4 plots the number of users’ reciprocal followers as a function of user ability. The number of reciprocal followers initially increases in user ability, but then decreases,²⁷ and eventually attains its minimum at high user abilities. In facet (e), we see clearly that the reciprocator-to-follower ratio is strictly decreasing in user ability. Low-ability users reciprocally follow the majority of their (few) followers, and users become more selective as their abilities increase. At the extreme, users of very high abilities follow back fewer than 0.5% of their followers. Facet (f) reports the fraction of a user’s followees that follow her back. About 25% of the followees of high-ability users follow them back, whereas the corresponding number is about 45% for medium-ability users. This relationship contrasts with the vertical model, in which all of a star’s followees follow them back. Allowing for some horizontal differentiation among the stars would be a natural way to capture the latter feature of the data.

Several predictions of our model are borne out in our #EconTwitter data. Middle-ability users are the ones who barter the most, with their attention bartering intensity increasing in their abilities, while low- and high-ability users barter the least—consistent with Proposition 3(iii).

5.6 Assortative matching

We find strong evidence of assortative matching. Figure 5 plots the ability ranges of users’ organic and reciprocal followees as a function of user ability. Panel 5a reports estimates from the #EconTwitter data. The left facet plots the distributions for users’ organic followees—followees who do not follow back—and the right facet plots the distribution for users’ reciprocal followees. In each facet, we report the median (black dots) and the interquartile range (pink-shaded area) of the ability distribution of the relevant category of followees, for each ability level. Panel 5b reports model predictions for the same quantities, reporting the median ability of users followees in the attention bartering equilibrium (black dots), and the ability range that users are willing to follow organically and reciprocally (pink-shaded area).


The data estimates and the predictions of the model are remarkably similar. Low-ability users are willing to reciprocate with any willing user who is not of extremely low quality. High-ability users form reciprocal relationships only with other high-ability users—note that these relationships are bidirectionally organic. The picture is completely different for organic followees: users follow unconditionally only high ability users, independent of their own abilities, both in our data and in our model.

Network formation models based on preferential attachment could explain some but not all of these patterns (Barabási and Albert, 1999). For example, highly right-skewed follower


²⁷Whereas stars have zero attention bartering relationships in the model, their gross quantity of bidirectional followers could be higher than for the highest-ability club, under a parameterization with more stars.

Figure 5: Ability ranges of users' followees

(a) #EconTwitter estimates



(b) Model predictions



Notes: This figure reports empirical estimates and theoretical predictions of the following behavior of users. Panel 5a plots the median and interquartile range of the distribution of followee abilities (y-axis) for #EconTwitter users belonging to the same ability percentile (x-axis). The left facet reports the distribution for organic followees—followees who do not follow the user, and the right facet for reciprocal followees—followees who are also followers. Panel 5b plots model predictions for the same quantities. The illustrated attention bartering equilibrium predictions are for the case of uniformly distributed types on $[-0.8, 0.2]$, $A(m) = 0.5\sqrt{m}$, and $c = 0.2$. See Section 5.3 for more details on the definition of abilities of #EconTwitter users. See Figure 2b for an alternative representation of the same attention bartering equilibrium.

distributions are consistent with rich-get-richer link formations, but would require more able users to have systematically joined the platform earlier. Had this been the case—and assuming that users do not “re-wire” frequently—it would also explain why higher ability users

tend to reciprocate with high ability users, as those are the users that were already present when they joined.²⁸

5.7 Alternative models of network formation

We examine whether other network formation models could explain the patterns in our data. We consider the preferential attachment model proposed by [Bollobás, Borgs, Chayes and Riordan \(2003\)](#) and the preferential attachment model with reciprocity proposed by [Cirkovic, Wang and Resnick \(2023\)](#). Neither model yields network statistics that match our empirical data. The full analyses can be found in [Appendix C](#).

6 Additional evidence for attention bartering

In [Section 5](#), we showed that attention bartering explains parsimoniously several striking patterns in our data. However, establishing directly that attention bartering exists is challenging: users do not sign an attention bartering “contract,” and other mechanisms could explain the relationship formation and the follow-for-follow patterns that we observe in our data. In this section, we provide additional evidence in support of attention bartering.

6.1 Direct evidence from surveys of Twitter and Instagram users

We conducted surveys with Twitter and Instagram users, asking them about their behavior related to attention bartering. Our goal is to provide direct evidence in support of attention bartering and to show that it generalizes to social media platforms beyond Twitter.


We used Prolific to recruit participants who were prescreened for Twitter and Instagram usage.²⁹ We created two distinct samples, comprising 496 participants in the Twitter survey, and 505 participants in the Instagram survey. [Figure 6](#) reports the key results of the two surveys. Each panel corresponds to a different question, the x-axis depicts the possible answers in increasing degree of agreement with the question, and the y-axis depicts the percentage of the users that selected each answer. The full survey and additional results can be found in [Appendix D](#).

We next describe the key results from the Twitter survey in what follows, but it is worth noting that the results of the Instagram user survey are strikingly similar. When asked the extent to which participants agree/disagree with the statement “All else equal, I am

²⁸We find very little relationship between user tenure and our ability measure; and our ability measure has substantially more predictive power than tenure does with respect to users’ follower-to-followee ratios. For more details, see [Appendix B](#).

²⁹Prolific is a leading online survey recruitment platform. When users join Prolific, they are asked to complete various “prescreeners” (questionnaires), ranging from basic demographics to more specific questions. We used the “social media” prescreener, which asks “Which of the following social media sites do you use on a regular basis (at least once a month)?”

Figure 6: A survey of Twitter and Instagram users on attention bartering.



Notes: This figure reports the key results from two surveys where we asked Twitter and Instagram users about their attention bartering behavior. The full survey and additional results can be found in Appendix D.

more likely to follow someone if they followed me,” 57.3% of participants said that they somewhat/strongly agree, 27.8% said that they somewhat/strongly disagree, and 14.9% were neutral (χ^2 test, p -value = 0). A higher proportion of participants agreed than disagreed (two-sided binomial test, p -value = 0).

Twitter users also exhibit reciprocal behavior in unfollowing other users: 61.3% of participants said that they somewhat or strongly agree with the statement “All else equal, I am more likely to unfollow someone if they unfollowed me,” 22.8 % said that they somewhat/strongly disagree, and 15.9% were neutral (χ^2 test, p -value = 0). A higher proportion of participants agreed than disagreed (two-sided binomial test, p -value = 0).

We also find strong evidence of reciprocity in how users engage with one another. When asked the extent to which they agree/disagree with the statement, “All else equal, I am more

likely to engage with tweets from someone that engages with my tweets,” 65.3% of participants said that they somewhat or strongly agree, 18.2 % said that they somewhat/strongly disagree, and 16.6% were neutral (χ^2 test, p -value = 0). A higher proportion of participants agreed than disagreed (two-sided binomial test, p -value = 0).

It is worth noting two additional results from our survey. First, we reframed the survey questions to reduce the “social desirability bias” by asking participants whether *people* are more likely to follow, engage, or unfollow reciprocally. The level of agreement with the reframed questions was higher than with the original, self-targeted questions: perhaps not surprisingly, people do not seem to like admitting that they follow, unfollow or engage with other users reciprocally. Second, to understand whether people monitor for reciprocal follows or unfollows, we asked participants what factor they typically notice when checking a user’s Twitter profile page. About 19.8% of participants said that they typically notice whether a user is following them when checking a user’s Twitter profile page.

6.2 The effect of reciprocity on relationship formation

If attention bartering exists, it should leave downstream empirical traces. One such trace is that reciprocal following should explain a large portion of relationship formation in our data. To test this hypothesis with our data, we evaluate how informative the fact that a user y follows a user x is in predicting whether user x follows user y . Our approach is to estimate two logistic regression specifications, with one omitting the information that user y follows user x .

We begin by describing our sample construction process. There are about 3 billion possible unilateral following relationships (directed edges) in the #EconTwitter data. To make estimating logistic regressions computationally feasible, we will sample 4 million edges among all possible edges—both edges that exist and that do not exist in the data. However, because the #EconTwitter graph is sparse, random sampling would yield a dependent variable distribution that is very biased towards the negative class (user x does not follow user y), and which could in turn bias our estimates.³⁰ To circumvent this “rare event” problem we will sample using the weighting scheme proposed by [King and Zeng \(2001\)](#), and sample without replacement an equal mix of existing edges and non-existing edges.³¹

Table 2 describes our independent variables and their summary statistics. It is worth noting the last three variables, which we construct to control for the similarity between each user x and user y , for all “user x follows user y ” edges in our sample. The similarity between

³⁰The network density of the #EconTwitter graph is about 0.004. In other words, only around 0.4% of the possible unilateral following relationships—the possible directed edges—exist in our data.

³¹This “equal shares sampling design” is close to the optimal design in our context ([Cosslett, 1981](#); [Imbens, 1992](#)). We weigh each observation i by $w_i = w_1 Y_i + w_0 (1 - Y_i)$, where $Y_i \in \{0, 1\}$ is the value of the target variable for that observation, τ is the fraction of 1’s in the population, \bar{y} is the fraction of 1’s in the sample, $w_1 = \frac{\tau}{\bar{y}}$, and $w_0 = \frac{1-\tau}{1-\bar{y}}$.

Table 2: Summary statistics of the independent variables.

Variable	Mean	Std. Dev.	Min	Max
Does x follow user y?	0.50	0.50	0.00	1
Does y follow user x?	0.20	0.40	0.00	1
User x’s number of followers	10157	23327	25	199962
User x’s number of followees	1595	1258	10	5000
Number of lists including user x	214	447	0.00	12077
User x’s number of tweets	11901	25409	50	1224343
User x’s number of likes	8533	22108	0.00	857558
User x’s ability	51	29	1	100
User y’s number of followers	17883	33556	25	199962
User y’s number of followees	1401	1160	10	5000
Number of lists including user y	378	702	0.00	12077
User y’s number of tweets	13779	27223	50	1224343
User y’s number of likes	8658	22758	0.00	857558
User y’s ability	59	29	1	100
Number of mutual followers	28	68	0.00	6371
Number of mutual followees	29	45	0.00	1033
Desc. Sim. between user x and user y	0.66	0.12	-0.10	1
Desc. sim. between x & y’s followers	0.79	0.09	0.07	0.96
Desc. sim. between x’s followees & y	0.79	0.09	0.01	0.96

user x and user y is the cosine similarity of the embeddings of their profile descriptions. The similarity between user x and user y ’s followers is the cosine similarity of the embedding of user x ’s profile description, and the average of the embeddings of user y ’s followers’ profile descriptions—removing user x , if user x belongs to that set (White, Togneri, Liu and Bennamoun, 2015; Coleman, 2020). Likewise, we calculate the similarity between user x ’s followees and user y .³² We standardize our continuous independent variables in order to be able to compare our estimates directly (Gelman, 2008).

Our full specification is:


$$y_i = \beta_0 + \beta_1 \text{FOLLOWS}_{y,x} + \gamma X + \epsilon_i$$

where y_i is an indicator variable for whether user x follows user y , $\text{FOLLOWS}_{y,x}$ is an indicator variable for whether user y follows user x , X is a vector of all other independent variables in Table 2, and ϵ_i is an error term. We estimate this specification excluding $\text{FOLLOWS}_{y,x}$ (“base model”) and then including it (“reciprocal model.”) Figure 7 reports the results of our analysis.

In the base model, the strongest predictors of whether user x follows user y are, as we

³²We use the transformer model BERT hosted by Hugging Face (Devlin, Chang, Lee and Toutanova, 2018).

Figure 7: Estimates of the link prediction model with and without reciprocation.



Notes: This figure reports normalized estimates from two logistic regressions. In both regressions, the dependent variable is an indicator variable for whether user x follows user y in our data. The “Base model” specification (red color) includes all independent variables reported on Table 2. The “Reciprocal model” specification (black color) additionally includes an indicator variable for whether user y follows user x , and its interaction with user x ’s ability. A 95% confidence interval is plotted around each estimate. The sizes of the “Reciprocal model” estimates are reported numerically. For more details on our the sample construction and the estimation strategy, see Section 6.2.

would expect, the numbers of their mutual followees and followers, as well as the ability of user y . In the reciprocal model, the magnitude of the first two estimates decreases substantially. Strikingly, we see that user x is about 70 times more likely to follow user y if user y follows them back. Note that we have added an interaction term between user x ’s ability and $\text{FOLLOWS}_{y,x}$, and that the estimate of this interaction term is substantial and negative. This estimate has a natural interpretation in light of our model: if user y follows user x and user x has high ability, it is less likely that the relationship is reciprocal. It is also worth noting that the estimated effect of user y ’s ability on user x following user y is higher in the reciprocal model. This result also has a natural interpretation in light of attention bartering: introducing the $\text{FOLLOWS}_{y,x}$ variable separates out the effect of following lower-ability users

just in order to be followed back, thereby increasing the ability estimate.

Our empirical exercise uses observational data, and hence we cannot rule out omitted variable bias completely. Nevertheless, we lay out three reasons why we believe our inferences are credible. First, the estimated effect of $\text{FOLLOWS}_{y,x}$ is so large that there would have to be a very significant predictor correlated with both $\text{FOLLOWS}_{x,y}$ and $\text{FOLLOWS}_{y,x}$ for the bias to remove the statistical robustness of the coefficient. Second, in the reciprocal model, the estimated effect of the number of mutual followers and mutual followees decreases dramatically, but little else shifts. This indicates that the $\text{FOLLOWS}_{y,x}$ is positively correlated with these measures. Third, the estimated effect of user y ’s ability increases in the reciprocal model, which is also consistent with attention bartering.

7 Discussion

We develop a model of social media where attention bartering shapes content production and consumption decisions—and hence network formation. Our model makes sharp predictions about users’ network statistics in equilibrium. We validate these predictions empirically using data from #EconTwitter, and show that other common network formation models fail to capture key empirical patterns. Furthermore, we show that attention bartering generalizes beyond #EconTwitter by conducting surveys with Twitter and Instagram users. Our findings suggest that (i) attention bartering takes place in social media platforms with very different designs and scopes, and (ii) attention bartering manifests both as users reciprocally following one other *and* as users reciprocally engaging with each others’ content.

In our model, attention bartering increases the amount of content produced and the utility of users who engage in attention bartering. If higher utilities imply more active users then platforms who monetize user activity—e.g., through advertising—may be better off allowing or even encouraging attention bartering.³³ On the other hand, attention bartering decreases the average quality of content produced. The reason is that users of lower abilities, who would have otherwise lurked, now form clubs and produce content.

Platform designers can encourage or dissuade attention bartering through monitoring costs. We observe substantial variation in these costs, both within and across social media platforms. For example, Twitter allows for attention bartering via reciprocal following by displaying prominently a “Follows you” tag. Similarly, Instagram shows which users have seen someone’s story, making attention bartering possible via reciprocal engagement. On the other hand, Instagram Reels and TikTok do not reveal the identity of the users who have watched a video—barring attention bartering via views—but display prominently the number of “likes” content has garnered.

³³[Srinivasan \(2023\)](#) provides empirical evidence on the attention-activity link. Interestingly, [Srinivasan](#) also shows that attention labor supply curve is concave—similar to our assumption in Section 3.

Another important implication is that users value receiving attention—even when they cannot monetize it. This can directly inform the design of features such as algorithmic recommendation systems. For example, [Srinivasan \(2023\)](#) formalizes this tradeoff in algorithm design: the platform may be better off showing some bad content, lowering the consumption utility of users who consume that content in order to retain the users who produce it. Attention bartering—and regulating it through monitoring costs—can be thought of as the market-based analogue of this algorithmic design.

Our work has limitations that follow up research could improve upon. On the modeling side, extensions to our model could consider dynamic elements of attention bartering including entry and exit, allow for substitution or complementarity in the types of users being followed, and introduce non-additive utilities and preferential attachment.³⁴ From an empirical perspective, a precise measurement of attention bartering would call for more fine-grained data such as content-level measures of engagement, more direct measures of ability, and longitudinal network formation data. Such data would allow researchers to conduct more direct tests of attention bartering, and estimate structurally its welfare effects.

8 Conclusion

This paper proposes a theory of attention bartering that can explain otherwise puzzling patterns of consumption and production on social media—a \$150bn industry and prominent information source. In our model, users of similar abilities form clubs and actively produce content, consuming the content of members of the same club—but they would have otherwise neither produced any content nor consumed each other’s content. Social media looks quite different from regular media in that most consumers are simultaneously content producers, a feature that our model accounts for with attention bartering.

We show empirically that several features of #EconTwitter, a salient Twitter community, are explicable through an economic, attention bartering lens. Our empirical findings are consistent with attention bartering taking place in #EconTwitter and our model predictions are largely borne out. Both in our model and in our data, medium-ability users are the keenest attention barterers, forming large fractions of their relationships reciprocally. On the other hand, low- and high-ability users follow others seemingly to only increase their consumption utilities, albeit for very different reasons: in our model, low-ability users would

³⁴For example, consider the case where every user may produce extremely valuable “star” content with a low probability that does not depend on her ability. Attention bartering then increases the amount of star content produced on the platform because it induces more users to produce content. Furthermore, by introducing additional edges to the network, attention bartering increases the probability that star content will become “viral.”

Conversely, if low-quality content spreads outside the clubs where it is produced—through algorithmic recommendations, new users who “join” these clubs unwittingly, or inter-club connections—existing users may receive lower total utility and dissuade would-be users from entering the platform.

like to barter but lack the opportunity to do so, whereas high-ability users have no incentive to barter. Twitter and Instagram user surveys corroborate our findings, and suggest that attention bartering takes place in social media platforms with different designs and scopes.

Attention bartering emerges because of the coexistence of production and consumption incentives for users. As such, it can occur even in the absence of algorithms steering users to form connections. In contrast to approaches that point to algorithms as the sole determinant of network structure, we offer an alternative incentive-based model of network formation, with algorithms playing no real role. With this simplification, our work complements existing work on social media that focuses on the role played by algorithms with relatively little attention paid to incentives. The degree to which the main patterns in the #EconTwitter data are explicable through a pure incentives perspective suggests that a sole focus on algorithmic explanations is, at the very least, incomplete.

References

- Alatas, Avi, Arun Chandrasekhar, Markus Mobius, Benjamin A. Olken, and Cindy Paladines**, “When celebrities speak: A nationwide twitter experiment promoting vaccination in indonesia,” *Working paper*, 2019.
- Allcott, Hunt and Matthew Gentzkow**, “Social media and fake news in the 2016 election,” *Journal of Economic Perspectives*, 2017, *31* (2), 211–36.
- , **Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow**, “The welfare effects of social media,” *American Economic Review*, 2020, *110* (3), 629–76.
- Bakshy, Eytan, Solomon Messing, and Lada A Adamic**, “Exposure to ideologically diverse news and opinion on Facebook,” *Science*, 2015, *348* (6239), 1130–1132.
- Barabási, Albert-László and Réka Albert**, “Emergence of scaling in random networks,” *Science*, 1999, *286* (5439), 509–512.
- Berman, Ron and Zsolt Katona**, “Curation algorithms and filter bubbles in social networks,” *Marketing Science*, 2020, *39* (2), 296–316.
- Bojanowski, Michał et al.**, “Do NBA teams avoid trading within their own division?,” *Network Science*, 2023, *11* (4), 657–669.
- Bollobás, Béla, Christian Borgs, Jennifer T Chayes, and Oliver Riordan**, “Directed scale-free graphs,” in “SODA,” Vol. 3 Baltimore, MD, United States 2003, pp. 132–139.
- Cirkovic, Daniel, Tiandong Wang, and Sidney I Resnick**, “Preferential attachment with reciprocity: properties and estimation,” *Journal of Complex Networks*, 2023, *11* (5), cnad031.
- Coleman, Benjamin**, “Why is it okay to average embeddings?,” <https://randorithms.com/2020/11/17/Adding-Embeddings.html> Nov 2020.
- Cosslett, Stephen**, “Efficient estimation of discrete choice models,” *Structural analysis of discrete data with econometric applications*, 1981.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova**, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” *CoRR*, 2018, *abs/1810.04805*.
- Gelman, Andrew**, “Scaling regression inputs by dividing by two standard deviations,” *Statistics in medicine*, 2008, *27* (15), 2865–2873.

- Imbens, Guido W**, “An efficient method of moments estimator for discrete choice models with choice-based sampling,” *Econometrica: Journal of the Econometric Society*, 1992, pp. 1187–1214.
- Jackson, Matthew O. and Asher Wolinsky**, “A strategic model of social and economic networks,” *Journal of Economic Theory*, 1996, 71 (1), 44–74.
- King, Gary and Langche Zeng**, “Logistic regression in rare events data,” *Political analysis*, 2001, 9 (2), 137–163.
- Krueger, Alan B**, “The economics of real superstars: The market for rock concerts in the material world,” *Journal of Labor Economics*, 2005, 23 (1), 1–30.
- Kwak, Haewoon, Changhyun Lee, Hosung Park, and Sue Moon**, “What is Twitter, a social network or a news media?,” in “Proceedings of the 19th international conference on World wide web” AcM 2010, pp. 591–600.
- Levy, Roe**, “Social media, news consumption, and polarization: Evidence from a field experiment,” *Working Paper*, 2020.
- Loeuille, Nicolas and Michel Loreau**, “Evolutionary emergence of size-structured food webs,” *Proceedings of the National Academy of Sciences*, 2005, 102 (16), 5761–5766.
- McNerney, James, Charles Savoie, Francesco Caravelli, Vasco M Carvalho, and J Doyne Farmer**, “How production networks amplify economic growth,” *Proceedings of the National Academy of Sciences*, 2022, 119 (1), e2106031118.
- Myers, Seth A, Aneesh Sharma, Pankaj Gupta, and Jimmy Lin**, “Information network or social network? The structure of the Twitter follow graph,” in “Proceedings of the 23rd International Conference on World Wide Web” 2014, pp. 493–498.
- Pennycook, Gordon, Ziv Epstein, Mohsen Mosleh, Antonio A Arechar, Dean Eckles, and David G Rand**, “Shifting attention to accuracy can reduce misinformation online,” *Nature*, 2021, pp. 1–6.
- Rosen, Sherwin**, “The economics of superstars,” *The American Economic Review*, 1981, 71 (5), 845–858.
- Sadri, Arif Mohaimin, Samiul Hasan, Satish V Ukkusuri, and Juan Esteban Suarez Lopez**, “Analysis of social interaction network properties and growth on Twitter,” *Social Network Analysis and Mining*, 2018, 8 (1), 1–13.
- Silverman, Bernard W**, *Density estimation for statistics and data analysis*, Vol. 26, CRC press, 1986.

Simon, Herbert, “Designing organizations for an information-rich world,” *Computers, Communications, and the Public Interest.*, 1971, pp. 37–72.

Srinivasan, Karthik, “Paying Attention,” 2023. Working Paper.

Toubia, Olivier and Andrew T Stephen, “Intrinsic vs. image-related utility in social media: Why do people contribute content to twitter?,” *Marketing Science*, 2013, *32* (3), 368–392.

Vicario, Michela Del, Gianna Vivaldo, Alessandro Bessi, Fabiana Zollo, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi, “Echo chambers: Emotional contagion and group polarization on Facebook,” *Scientific Reports*, 2016, *6* (1), 1–12.

White, Lyndon, Roberto Togneri, Wei Liu, and Mohammed Bennamoun, “How well sentence embeddings capture meaning,” in “Proceedings of the 20th Australasian document computing symposium” 2015, pp. 1–8.

A Proofs for theoretical results

Proof of Theorem 1. The first condition in the proposition's statement is exactly the first condition in Definition 1. So we need to show, assuming this condition on f_1 holds, that the second condition in the proposition's statement is equivalent to the second condition in Definition 1. Henceforth, fix $(f_1, f_2) \in \mathcal{F}_1 \times \mathcal{F}_2$ satisfying the first condition.

We begin with some payoff computations. First, observe, any $\theta, \hat{\theta} \in \Theta$ have

$$\begin{aligned} & [1 - f_1(\hat{\theta}, \theta)]u(\hat{\theta}, \theta) - c(\hat{\theta}) + [1 - f_1(\hat{\theta}, \theta)]b(\hat{\theta})A'(m(\hat{\theta}|f_1, f_2)) \\ = & \min\{0, u(\hat{\theta}, \theta)\} - c(\hat{\theta}) + [1 - f_1(\hat{\theta}, \theta)]b(\hat{\theta})A'(m(\hat{\theta})), \end{aligned}$$

which is equal, for any $\theta \in \Theta$ and almost every $\hat{\theta} \in \Theta$, to

$$\varphi(\theta, \hat{\theta}) := \min\{0, u(\theta, \hat{\theta})\} - c(\theta) + b(\theta)A'(m(\theta))\mathbf{1}_{u(\hat{\theta}, \theta) \leq 0}.$$

Define now, for each $\theta \in \Theta$, the set \mathcal{F}_2^θ of all $\tilde{f}_2 \in \mathcal{F}_2$ such that

- Every $\hat{\theta}, \tilde{\theta} \in \Theta \setminus \{\theta\}$ have $\tilde{f}_2(\hat{\theta}, \tilde{\theta}) = f_2(\hat{\theta}, \tilde{\theta})$;
- Every $\hat{\theta} \in \Theta$ with $\tilde{f}_2(\theta, \hat{\theta}) > f_2(\theta, \hat{\theta})$ has $\varphi(\hat{\theta}, \theta) > 0$.

The above calculation tells us (f_1, f_2) is an equilibrium if and only if every $\theta \in \Theta$ has $V(\theta|f_1, f_2) \geq V(\theta|f_1, \tilde{f}_2)$ for each $\tilde{f}_2 \in \mathcal{F}_2^\theta$.

Next observe that, for each $\theta \in \Theta$, the function $V(\theta|f_1, \cdot)$ is concave, and the set \mathcal{F}_2^θ is star-shaped at f_2 . Hence, the function is maximized at f_2 if and only if it is locally maximized there in each direction. Given some $\tilde{f}_2 \in \mathcal{F}_2$, each $\epsilon \in [0, 1]$ has, letting $f_2 + \epsilon(\tilde{f}_2 - f_2)$,

$$\begin{aligned} V(\theta|f_1, f_2^\epsilon) &= \int_{\Theta} \left[f_2^\epsilon(\theta, \cdot) [u(\theta, \cdot) - c(\theta)] + [1 - f_2^\epsilon(\theta, \cdot)] f_1(\theta, \cdot) u(\theta, \cdot) \right] d\mu \\ &\quad + b(\theta)A(m(\theta|f_1, f_2^\epsilon)) \\ &= \int_{\Theta} \left[f_2^\epsilon(\theta, \cdot) [u(\theta, \cdot) - c(\theta)] + [1 - f_2^\epsilon(\theta, \cdot)] \max\{0, u(\theta, \cdot)\} \right] d\mu \\ &\quad + b(\theta)A(m(\theta|f_1, f_2^\epsilon)). \end{aligned}$$

Hence,

$$\begin{aligned}
\left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} V(\theta|f_1, f_2^\epsilon) &= \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} \int_{\Theta} f_2^\epsilon(\theta, \cdot) [u(\theta, \cdot) - c(\theta) - \max\{0, u(\theta, \cdot)\}] d\mu \\
&\quad + b(\theta) \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} A(m(\theta|f_1, f_2^\epsilon)) \\
&= \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} \int_{\Theta} f_2^\epsilon(\theta, \cdot) [\min\{0, u(\theta, \cdot)\} - c(\theta)] d\mu \\
&\quad + b(\theta) A'(m(\theta)) \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} m(\theta|f_1, f_2^\epsilon) \\
&= \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} \int_{\Theta} f_2^\epsilon(\theta, \cdot) [\min\{0, u(\theta, \cdot)\} - c(\theta)] d\mu \\
&\quad + b(\theta) A'(m(\theta)) \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} \int_{\Theta} \{f_2^\epsilon(\theta, \cdot) + [1 - f_2^\epsilon(\theta, \cdot)]f_1(\cdot, \theta)\} d\mu \\
&= \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} \int_{\Theta} f_2^\epsilon(\theta, \cdot) [\min\{0, u(\theta, \cdot)\} - c(\theta)] d\mu \\
&\quad + b(\theta) A'(m(\theta)) \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} \int_{\Theta} f_2^\epsilon(\theta, \cdot) [1 - f_1(\cdot, \theta)] d\mu \\
&= \left. \frac{\partial}{\partial \epsilon} \right|_{\epsilon=0} \int_{\Theta} f_2^\epsilon(\theta, \cdot) \varphi(\theta, \cdot) d\mu \\
&= \int_{\Theta} [\tilde{f}_2(\theta, \cdot) - f_2(\theta, \cdot)] \varphi(\theta, \cdot) d\mu.
\end{aligned}$$

Given the above observations and calculations, (f_1, f_2) is an equilibrium if and only if every $\theta \in \Theta$ and $\tilde{f}_2 \in \mathcal{F}_2^\theta$ have

$$\int_{\Theta} [\tilde{f}_2(\theta, \cdot) - f_2(\theta, \cdot)] \varphi(\theta, \cdot) d\mu \leq 0.$$

But this condition holds if every $\theta \in \Theta$ and almost every $\hat{\theta} \in \Theta$ have

$$f_2(\theta, \hat{\theta}) = \begin{cases} 1 & : \varphi(\theta, \hat{\theta}), \varphi(\hat{\theta}, \theta) > 0 \\ 0 & : \varphi(\theta, \hat{\theta}) < 0; \end{cases}$$

and otherwise does not hold as witnessed by $\tilde{f}_2 \in \mathcal{F}_2^\theta$ with

$$\tilde{f}_2(\tilde{\theta}, \hat{\theta}) := \begin{cases} 1 & : \theta \in \{\tilde{\theta}, \hat{\theta}\} \text{ and } \varphi(\tilde{\theta}, \hat{\theta}), \varphi(\hat{\theta}, \tilde{\theta}) > 0 \\ 0 & : \tilde{\theta} = \theta \text{ and } \varphi(\theta, \hat{\theta}) < 0 \\ 0 & : \hat{\theta} = \theta \text{ and } \varphi(\theta, \tilde{\theta}) < 0 \\ f_2(\tilde{\theta}, \hat{\theta}) & : \text{otherwise.} \end{cases}$$

Now, because $\frac{1}{b(\theta)} \varphi(\theta, \hat{\theta}) = A'(m(\theta)) \mathbf{1}_{u(\hat{\theta}, \theta) \leq 0} - \kappa(\theta, \hat{\theta})$ for every $\theta, \hat{\theta} \in \Theta$, our arguments

to this point tell us (f_1, f_2) is an equilibrium if and only if every $\theta \in \Theta$ and almost every $\hat{\theta} \in \Theta$ have

$$f_2(\theta, \hat{\theta}) = \begin{cases} 1 & : \kappa(\theta, \hat{\theta}) < A'(m(\theta))\mathbf{1}_{u(\hat{\theta}, \theta) \leq 0} \text{ and } \kappa(\hat{\theta}, \theta) < A'(m(\hat{\theta}))\mathbf{1}_{u(\theta, \hat{\theta}) \leq 0} \\ 0 & : \kappa(\theta, \hat{\theta}) > A'(m(\theta))\mathbf{1}_{u(\hat{\theta}, \theta) \leq 0}. \end{cases}$$

To complete the proof, note that the augmented cost κ is always strictly positive, and so the latter equation for f_2 is equivalent to the proposition's second condition. \square

Proof of Proposition 1. Let (f_1, f_2) be some equilibrium, and let $m(\cdot) := m(\cdot | f_1, f_2)$. Assume for a contradiction that $m(\theta) < m(\tilde{\theta})$. Then, almost every type $\hat{\theta} \in \Theta$ has:

- $u(\theta, \hat{\theta}) > 0$ if $u(\tilde{\theta}, \hat{\theta}) \geq 0$;
- $\kappa(\theta, \hat{\theta}) \leq A'(m(\tilde{\theta})) < A'(m(\theta))$ if $\kappa(\tilde{\theta}, \hat{\theta}) \leq A'(m(\tilde{\theta}))$;
- $\kappa(\hat{\theta}, \theta) < A'(m(\hat{\theta}))$ if $\kappa(\hat{\theta}, \tilde{\theta}) \leq A'(m(\hat{\theta}))$.

Hence, by Theorem 1, we have $f_1(\hat{\theta}, \theta) \geq f_1(\hat{\theta}, \tilde{\theta})$ and $f_2(\hat{\theta}, \theta) \geq f_2(\hat{\theta}, \tilde{\theta})$ for almost every $\hat{\theta} \in \Theta$, contradicting the hypothesis that $m(\theta) < m(\tilde{\theta})$. \square

Proof of Claim 1. First, because $A : [0, 1] \rightarrow \mathbb{R}$ is concave, and is differentiable on $(0, 1]$ and globally continuous, it follows that $A'(0) \in [0, \infty]$ exists too, and the function $A' : [0, 1] \rightarrow \mathbb{R}_+ \cup \{\infty\}$ is continuous and strictly decreasing.

Now let us show, by induction on $n \in \mathbb{N}$, that $\theta_n \geq c - A'(0)$, and the equation $A'(G(\theta_{n-1}) - G(\theta_n)) = c - \theta_n$ has a unique solution $\theta_n \in [-\infty, \theta_{n-1}]$. Suppose $n \in \mathbb{N}$ and these properties hold for $n - 1$ if $n > 1$. Clearly, $\theta_n = -\infty$ is as desired if $\theta_{n-1} = -\infty$, so focus on the case that $\theta_{n-1} \in \mathbb{R}$. If $A'(G(\theta_{n-1})) = \infty$ —that is, if $G(\theta_{n-1}) = 0$ and $A'(0) = \infty$ —then every $\theta \in (-\infty, \theta_{n-1}]$ has $A'(G(\theta_{n-1}) - G(\theta)) = \infty > c - \theta$, and so $\theta_n = -\infty$ is as desired. So let us further focus on the case that $A'(G(\theta_{n-1})) < \infty$.

In this case, define the continuous function

$$\begin{aligned} \lambda_n : (-\infty, \theta_{n-1}] &\rightarrow \mathbb{R} \cup \{\infty\} \\ \theta_n &\mapsto A'(G(\theta_{n-1}) - G(\theta_n)) - (c - \theta_n). \end{aligned}$$

Being strictly increasing, λ_n has at most one zero. Moreover, $\lambda_n(\theta_{n-1}) \geq 0$ by the inductive hypothesis, and $\lambda_n(\theta_n)$ is negative for low enough θ_n in its domain. Hence the intermediate value theorem delivers a root. Moreover, in this case, $c - (-\infty) < A'(G(\theta_{n-1}))$, so that this θ_n uniquely solves the equation, completing the inductive step.

Finally, we characterize the limit. Because $(G(\theta_{n-1}) - G(\theta_n))_{n=1}^{\infty}$ is a summable series, it follows that $G(\theta_{n-1}) - G(\theta_n) \rightarrow 0$ as $n \rightarrow \infty$. Hence,

$$\theta_{\infty} = \lim_{n \rightarrow \infty} \theta_n = \lim_{n \rightarrow \infty} A'(G(\theta_{n-1}) - G(\theta_n)) - c = A'(0) - c,$$

as required. \square

Proof of Proposition 2. Observe, any $\hat{\theta} \in \Theta$ with $\hat{\theta} < 0$ has $\kappa(\cdot, \hat{\theta}) = c - \hat{\theta}$. Hence, Theorem 1 tells us a pair $(f_1, f_2) \in \mathcal{F}_1 \times \mathcal{F}_2$, inducing follower count $m(\cdot) = m(\cdot | f_1, f_2)$, is an equilibrium if and only if:

1. Every $\theta, \hat{\theta} \in \Theta$ have $f_1(\theta, \hat{\theta}) = \mathbf{1}_{\hat{\theta} > 0}$.
2. Every $\theta \in \Theta$ and almost every $\hat{\theta} \in \Theta$ have

$$f_2(\theta, \hat{\theta}) = \begin{cases} 1 & : \hat{\theta} < 0, \theta < 0, c - \hat{\theta} < A'(m(\theta)), \text{ and } c - \theta < A'(m(\hat{\theta})) \\ 0 & : \theta > 0 \text{ or } c - \hat{\theta} > A'(m(\theta)). \end{cases}$$

It follows directly that (f_1^*, f_2^*) is an equilibrium. In particular, types in $\hat{\theta} \in (\theta_n, \theta_{n-1}]$ can only find types in $(-\infty, \theta_{n-1}]$ as willing partners, but only want to barter with types in (θ_n, ∞) .

Now, consider an arbitrary equilibrium $(f_1, f_2) \in \mathcal{F}_1 \times \mathcal{F}_2$, and let $m(\cdot) := m(\cdot | f_1, f_2)$. We know $f_1 = f_1^*$ from above; we now want to show f_2 coincides almost everywhere with f_2^* . To begin with, Corollary 3 tells us any $\theta \in \Theta$ with $\theta > 0$ has $f_2(\theta, \cdot) = 0 = f_2^*(\theta, \cdot)$ almost everywhere.

To prove f_2 agrees with f_2^* almost everywhere, we separately consider two cases. First, consider the case that $A'(0) \leq c$. As noted above, θ has zero mass of barterers among types $\hat{\theta} > 0$. Moreover, any type $\hat{\theta} < 0$ has $c - \hat{\theta} > c \geq A'(0) \geq A'(m(\theta))$; hence, θ attention barterers with zero mass of types $\hat{\theta} < 0$.

Now, we turn to the case that $A'(0) \geq c$. Let $\{\theta_n\}_{n=1}^{\infty}$ be as defined in Claim 1, and let $\Theta_n := \Theta \cap (\theta_n, \theta_{n-1})$ for each $n \in \mathbb{N}$. Let us show by induction that, for every $n \in \mathbb{N}$, every $\theta \in \Theta_n$ has $f_2(\theta, \cdot) =_{\text{a.e.}} \mathbf{1}_{\Theta_n}$, which will deliver the proposition. To that end, take some $n \in \mathbb{N}$, and suppose the property holds for every lower index. We have nothing to show if $\Theta_n = \emptyset$, so assume without loss that Θ_n is nonempty. Defining the function

$$\begin{aligned} \tau : \Theta &\rightarrow \mathbb{R} \cup \{-\infty\} \\ \theta &\mapsto c - A'(m(\theta)), \end{aligned}$$

Theorem 1 tells us every $\theta \in \Theta$ and almost every $\hat{\theta} \in \Theta$ with $\theta, \hat{\theta} < 0$ have $f_2(\theta, \hat{\theta}) =$

$\mathbf{1}_{\theta > \tau(\hat{\theta})}$ and $\hat{\theta} > \tau(\theta)$. In what follows, we derive properties of τ and infer properties of the equilibrium. Throughout the argument we make use of the fact that m is nondecreasing (by Proposition 1), and hence τ is as well.

Let us first show that $\check{\theta}_n := \sup\{\tau(\theta) : \theta \in \Theta_n\} \leq \theta_{n-1}$. Assume otherwise for a contradiction. In this case, given that τ is monotone, some $\epsilon > 0$ exists such that $(\theta_{n-1} - \epsilon, \theta_{n-1}) \cap \Theta$ is nonempty and every θ in this set has $\tau(\theta) > \theta_{n-1}$. Hence, any $\theta \in (\theta_{n-1} - \epsilon, \theta_{n-1}) \cap \Theta$ has $f_2(\theta, \hat{\theta}) = 0$ for almost every $\hat{\theta} \in \Theta$ with $\hat{\theta} < \theta_{n-1}$, and so (by the inductive hypothesis) has $m(\theta) = 0$. Hence, monotonicity of τ implies $\check{\theta}_n = c - A'(0) \leq \theta_{n-1}$, a contradiction.

Now, let us see that $\check{\Theta}_n := \Theta \cap (\check{\theta}_n, \theta_{n-1})$ essentially forms a clique. Hence, defining $\check{\Theta}_n := \Theta \cap (\check{\theta}_n, \theta_{n-1})$, every $\theta, \hat{\theta} \in \check{\Theta}_n$ necessarily have $\tau(\hat{\theta}) \leq \check{\theta}_n < \theta_n$, and symmetrically have $\tau(\hat{\theta}) < \theta_n$. It follows that every $\theta \in \check{\Theta}$ and almost every $\hat{\theta} \in \check{\Theta}$ have $f_2(\theta, \hat{\theta}) = 1$.

Defining $m^* := G(\theta_{n-1}) - G(\check{\theta}_n)$, observe that the inductive step (hence the proposition) will follow if we can show $m(\theta) = m^*$ for every $\theta \in \check{\Theta}_n$. Indeed, in this case, every $\theta \in \check{\Theta}_n$ has $f_2(\theta, \cdot) \geq_{\text{a.e.}} \mathbf{1}_{\check{\Theta}_n}$ and $\int f_2(\theta, \cdot) d\mu = m^* = \int \mathbf{1}_{\check{\Theta}_n} d\mu$, and hence $f_2(\theta, \cdot) =_{\text{a.e.}} \mathbf{1}_{\check{\Theta}_n}$. Moreover, in this case, $\check{\theta}_n = \tau(m^*) = c - A'(G(\theta_{n-1}) - G(\check{\theta}_n))$, so that $\check{\theta}_n = \theta_n$.

So all that remains is to establish $m(\theta) = m^*$ for every $\theta \in \check{\Theta}_n$. To that end, note every $\theta \in \check{\Theta}_n$ has $\mathbf{1}_{\check{\Theta}_n} \leq_{\text{a.e.}} f_2(\theta, \cdot) \leq_{\text{a.e.}} \mathbf{1}_{\check{\Theta}_n \cup [\Theta \cap (\tau(\theta), \check{\theta}_n)]}$. Hence,

$$m(\theta) = \int_{\Theta} f_2(\theta, \cdot) d\mu \in [m^*, m^* + G(\check{\theta}_n) - G(\tau(\theta))].$$

In particular, $m|_{\check{\Theta}_n} \geq m^*$. Therefore, letting $\{\check{\theta}_{n,k}\}_{k=1}^{\infty} \subseteq \check{\Theta}_n$ be some sequence converging to $\sup \check{\Theta}_n$, monotonicity of m will deliver $m|_{\check{\Theta}_n} = m^*$ if we can show $\lim_{k \rightarrow \infty} m(\check{\theta}_{n,k}) = m^*$. Given that $m(\check{\theta}_{n,k}) - m^* \in [0, G(\check{\theta}_n) - G(\tau(\check{\theta}_{n,k}))]$ for $k \in \mathbb{N}$ and G is continuous, it suffices to see that $\tau(\check{\theta}_{n,k}) \rightarrow \check{\theta}_n$ as $k \rightarrow \infty$. But the latter holds by monotonicity of τ . The proposition follows. \square

Proof of Proposition 3. Proposition 2 tells us that lurkers are those users with ability below θ_{∞} , that stars are those users with ability above θ_0 , and that for each $n \in \mathbb{N}$ the users in the n^{th} club $(\theta_n, \theta_{n-1}]$ are followed by the users in their own club and nobody else, and that their followees consist of their followers and the stars.

Given these observations, everything in the proposition other than the final sentence of part (ii) is satisfied if higher-ability users have a weakly higher number of followers—a feature that Proposition 1 guarantees.

All that remains is to show, given $A'(G(0)) + \underline{\theta} < c < A'(0)$, that the number of followees is non-monotone in a user's ability. Because stars have zero followees and we have argued the number of followees of a non-star is weakly increasing in her ability, all we need to verify

is that the number of followees is not the same for every non-star. To that end, observe

$$A'(G(0)) = A'(G(\theta_0) - G(\underline{\theta})) < c - \underline{\theta}$$

and

$$A'(0) = A'(G(\theta_0) - G(\theta_0)) > c = c - \theta_0$$

Hence, that $\theta \mapsto A'(G(\theta_0) - G(\underline{\theta})) - [c - \theta]$ is increasing implies $\underline{\theta} < \theta_1 < \theta_0$. That is, the first club $(\theta_1, \theta_0]$ constitutes a nontrivial fraction of the non-stars. That is, it is not the case that every non-star user belongs to a single club. Finally, that

$$\theta_1 = c - A'(G(0) - G(\theta_1)) > c - A'(0) = \theta_\infty = \lim_{n \rightarrow \infty} \theta_n$$

implies $\theta_2 < \theta_1$, and so non-stars outside the first club have *strictly* fewer followers than those in the first club. \square


B Additional empirical results

B.1 Follower, followee, and ratio distributions of all users

Figure 8 reports the distributions of followers, followees, and follower-to-followee ratios of users in the entire sample. For each outcome, we report a Gaussian kernel density estimate of its distribution with the bandwidth selected using Silverman’s rule of thumb (Silverman, 1986), as well as the estimated median (solid red vertical line) and the mean (dashed red vertical line).

Compared to the average #EconTwitter user, Twitter users in our data have on average fewer followers and followees, and lower ratios: the median user in the entire data set has 266 followers, 524 followees, and a ratio equal to 0.52. Only 25.1% of the users have ratio higher than one in the entire sample, and the distribution of ratios exhibits a noticeable kink at the unit ratio—this is consistent with Proposition 3(i). It is worth noting that the effects of Twitter’s followee limits are easily discernible in the kinks of the middle facet.


Figure 8: Distributions of followers, followees, and ratios of all users



Notes: This figure reports kernel density estimates of the probability density function of the distributions of followers, followees, and follower-to-followee ratios in the entire data. For each facet, the vertical red lines depict the median (solid) and the mean (dashed) of the corresponding distribution, and the kernel bandwidth is selected using Silverman’s rule of thumb (Silverman, 1986). See Figure 3 for density estimates of the same statistics for #EconTwitter users.

B.2 More details on the empirical estimates of Figure 4

Figure 9: Unscaled #EconTwitter estimates of network statistics.



Notes: This figure reports unscaled #EconTwitter estimates of various network statistics, as a function of user ability. For more details on the construction of the estimates, see the description of Figure 4.

B.3 Ability is not highly correlated with tenure

One alternative explanation for the patterns we observe in the #EconTwitter data can be found with preferential attachment. Namely, if more able users join the platform earlier, then these users would attract more followers—ability is not the reason for attracting followers, but time of entry is. Furthermore, these users may also follow more able users on average because those users were active at the time of relationship formation.

Table 3 examines the relation of tenure, ability, and users’ follower-to-followee ratios. Tenure is positively correlated with ability and with follower-to-followee ratios, but it only explains about 4% of the variability in both quantities. In contrast, user ability is about eight times more predictive of a users’ follower-to-followee ratio. The lack of predictive power of user tenure suggests that the preferential attachment explanation is not strongly supported in our data.

Table 3: The relation of tenure, ability, and follower-to-followee ratio.

	<i>Dependent variable:</i>		
	Ability (percentile)	Follower/followee Ratio	
	(1)	(2)	(3)
Tenure (days)	0.006*** (0.0001)	0.0003*** (0.00001)	
Ability (percentile)			0.025*** (0.0001)
Constant	33.463*** (0.375)	0.777*** (0.016)	0.267*** (0.008)
Observations	55,230	55,230	55,230
R ²	0.041	0.042	0.344
Adjusted R ²	0.041	0.042	0.344
Residual Std. Error (df = 55228)	28.271	1.184	0.980
F Statistic (df = 1; 55228)	2,337.627***	2,436.300***	28,907.110***

Notes: This table reports regressions where the dependent variable is ability (column 1) and follower-to-followee ratio (columns 2 and 3). The independent variables are users’ tenure on the platform, and users’ abilities. For more details on the definition of user ability, see Section 5.3 Significance indicators: $p \leq 0.1$: ‡, $p \leq 0.05$: *, $p \leq 0.01$: **, and $p \leq .001$: ***.

C Evaluating alternative models of network formation

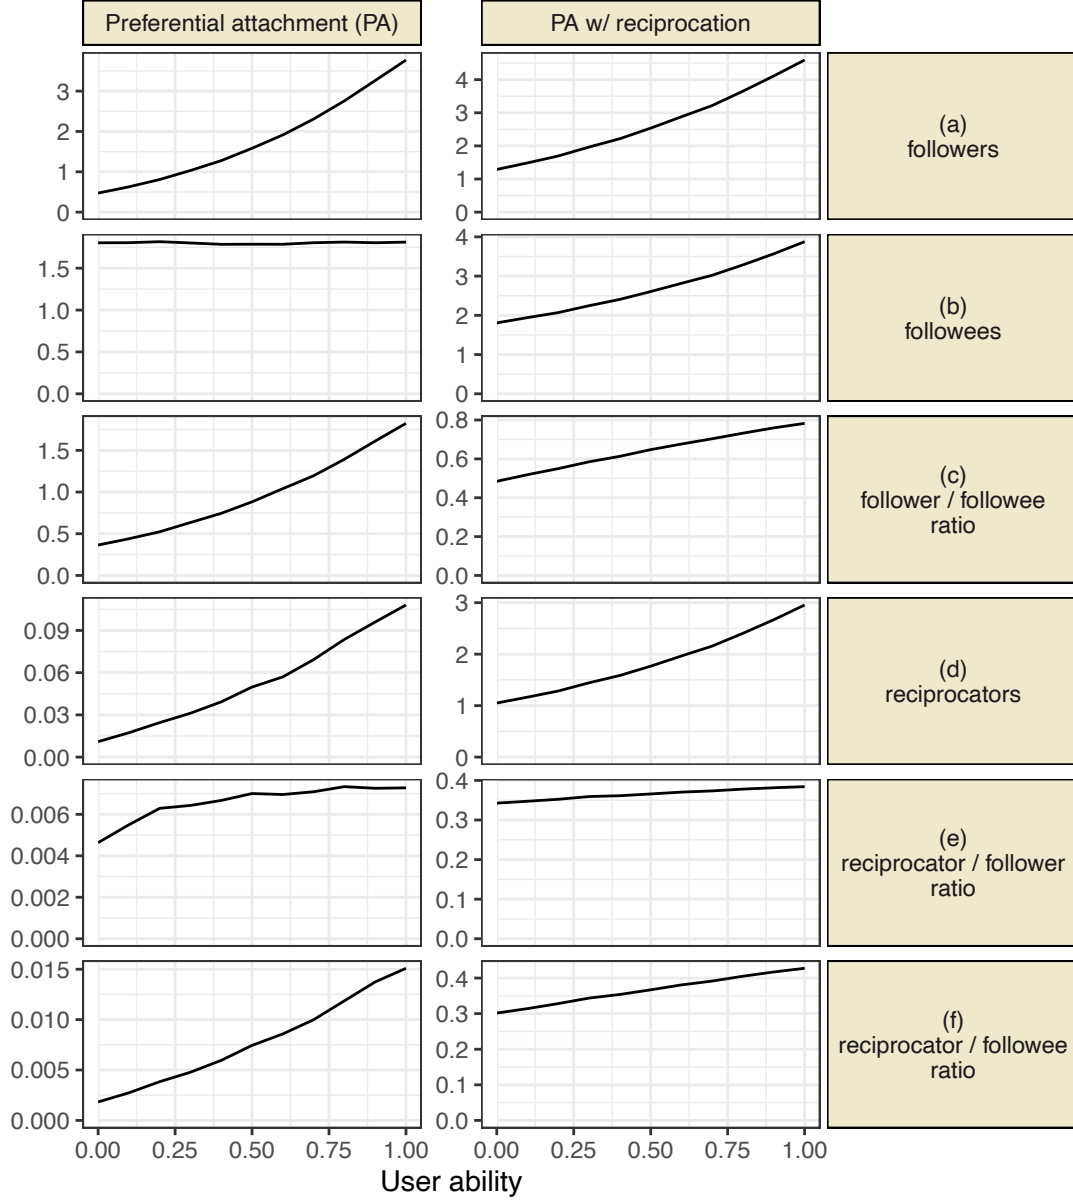
One criticism of our empirical approach is that other network formation models could explain the patterns in our data. We first consider the seminal preferential attachment model, extended to directed graphs (Bollobás et al., 2003). In this model, each node is characterized by its in-degree, out-degree, “in-fitness” (similar to ability in our model), and “out-fitness” (a user’s propensity to follow other users). Edges are added probabilistically based on the characteristics of each node. An incoming edge to a node is added with probability proportional to the in-degree (i.e., popularity) and in-fitness (i.e., “ability”) of the node, and an outgoing edge from a node is added with probability proportional to the out-fitness of the node. Importantly, there is no concept of reciprocity in this model.

We perform network formation simulations using the preferential attachment model. The left panel in Figure 10 plots the results for the case where the parameters of the preferential attachment model are: $\alpha = 0.3$, $\beta = 0.5$, $\gamma = 0.2$, $\delta_{in} = 1$, $\delta_{out} = 1$. The preferential attachment model captures correctly the relationship of a user’s followers and follower-followee ratio as a function of that user’s ability. However, it fails to predict several other patterns: in contrast to the #EconTwitter data, the number of followees is constant in ability, and the the reciprocator-to-follower ratio is increasing in ability. Conducting the same simulations using different parameters yielded qualitatively similar results. This suggests that the preferential attachment model does not capture users’ following choices on EconTwitter.

Next, we extend the preferential attachment model with reciprocation as proposed by Cirkovic et al. (2023). This model extends the directed preferential attachment model by allowing a reciprocal relationship to be formed with a fixed probability ρ every time an edge is added to the network. The right panel in Figure 10 plots the results for the case where the parameters of the extended model are: $\alpha = 0.3$, $\beta = 0.5$, $\gamma = 0.2$, $\delta_{in} = 1$, $\delta_{out} = 1$, $\rho = 0.5$. This model still does not match all the patterns in the #EconTwitter data. Notably, although the number of followees is increasing in ability due to reciprocity, it does not have the characteristic dip that is predicted by our model and observed in the #EconTwitter data.

Recent work has attempted to introduce further extensions to network formation models to allow them to better match empirical patterns for specific contexts. For example, McNerney, Savoie, Caravelli, Carvalho and Farmer (2022) focus on supply networks and use a centrality measure to predict price changes from firms’ productivity improvements; Bojanowski et al. (2023) focus on sports trading and use a weighted exponential random graph model to predict team trading behavior; and Loeuille and Loreau (2005) focus on predator-prey networks and use model the evolution of this network, allowing for competition between species. Our approach is more similar in spirit to these models: we focus specifically on the role of attention in the production and consumption of content in social networks.

Figure 10: Predictions of alternate network formation models



Notes: This figure reports the predictions of networks statistics using alternate network formation models. The left panel plots the predictions of the preferential attachment model (Bollobás et al., 2003) with parameters: $\alpha = 0.3$, $\beta = 0.5$, $\gamma = 0.2$, $\delta_{in} = 1$, $\delta_{out} = 1$. The right panel plots the predictions of the extended preferential attachment model with reciprocation (Cirkovic et al., 2023) with an additional reciprocation parameter $\rho = 0.5$.

D Additional details on the Twitter and Instagram surveys

Participants were recruited on Proffic. In total, we received responses from 496 participants for the Twitter survey, and 505 for the Instagram survey. The effective average hourly compensation rate was \$21 per hour for the Twitter survey, and \$14 per hour for the Instagram survey. Tables 4 and 5 report the questions and results from the Twitter and Instagram surveys, respectively.

Table 4: Twitter survey questions

(a) Twitter usage		Proportion
How long have you had a Twitter account?	Less than 6 months	0.010
	6 months to 1 year	0.018
	1 to 2 years	0.083
	2 to 5 years	0.304
	More than 5 years	0.585
How often do you use Twitter?	Once a month or less	0.043
	A few times a week	0.222
	Once a day	0.116
	Multiple times a day	0.619
How many Twitter followers do you have?	Less than 50	0.383
	50 to 500	0.395
	501 to 2,000	0.135
	2,001 to 10,000	0.077
	More than 10,000	0.010
How many users do you follow on Twitter?	Less than 50	0.226
	50 to 500	0.484
	501 to 1,000	0.159
	1,001 to 5,000	0.113
	More than 5,000	0.018

(b) Reciprocity behavior

	Strongly disagree	Somewhat disagree	Neutral	Somewhat agree	Strongly agree
<i>To what extent do you agree/disagree with the following statement:</i>					
“All else being equal, I’m more likely to engage with tweets (i.e., like, comment, retweet) from someone that engages with my tweets”	0.059	0.123	0.166	0.331	0.321
“All else being equal, people are more likely to engage with tweets (i.e., like, comment, retweet) from users who engage with their tweets”	0.015	0.066	0.164	0.424	0.330
“All else being equal, I’m more likely to follow someone if they followed me first”	0.113	0.165	0.149	0.329	0.244
“All else being equal, people are more likely to follow users who followed them first”	0.011	0.070	0.149	0.473	0.297
“All else being equal, I’m more likely to unfollow someone if they unfollowed me first”	0.077	0.151	0.159	0.258	0.355
“All else being equal, people are more likely to unfollow users who unfollowed them first”	0.008	0.072	0.121	0.406	0.393

(c) Profile checking behavior

	Proportion
<i>When you check someone’s Twitter profile page, what do you typically pay attention to? (select all that apply)</i>	
Their recent tweets	0.908
Their profile description	0.635
Their profile photo	0.588
Their follower count	0.286
Whether they are following you	0.198

Table 5: Instagram survey questions

(a) Instagram usage		Proportion
How long have you had an Instagram account?	Less than 6 months	0.004
	6 months to 1 year	0.010
	1 to 2 years	0.055
	2 to 5 years	0.220
How often do you use Instagram?	More than 5 years	0.711
	Once a month or less	0.031
	A few times a week	0.122
	Once a day	0.132
How many Instagram followers do you have?	Multiple times a day	0.715
	Less than 50	0.168
	50 to 500	0.515
	501 to 2,000	0.253
How many users do you follow on Instagram?	2,001 to 10,000	0.059
	More than 10,000	0.004
	Less than 50	0.131
	50 to 500	0.509
	501 to 1,000	0.240
	1,001 to 5,000	0.099
	More than 5,000	0.022

(b) Reciprocity behavior

	Strongly disagree	Somewhat disagree	Neutral	Somewhat agree	Strongly agree
<i>To what extent do you agree/disagree with the following statement:</i>					
“All else being equal, I’m more likely to engage with posts (i.e., like, comment) from someone that engages with my posts”	0.060	0.098	0.131	0.392	0.319
“All else being equal, people are more likely to engage with posts (i.e., like, comment) from users who engage with their posts”	0.026	0.044	0.135	0.495	0.300
“All else being equal, I’m more likely to follow someone if they followed me first”	0.079	0.141	0.147	0.386	0.248
“All else being equal, people are more likely to follow users who followed them first”	0.014	0.079	0.125	0.515	0.267
“All else being equal, I’m more likely to unfollow someone if they unfollowed me first”	0.057	0.091	0.156	0.250	0.446
“All else being equal, people are more likely to unfollow users who unfollowed them first”	0.010	0.034	0.105	0.402	0.450