Appendix

Apostolos Filippas
Fordham

John J. Horton
MIT & NBER

Elliot Lipnowski
Yale

Prasanna Parasurama
Emory

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 $\mathcal{F}_2^\theta$ 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 $\mathcal{F}_2^\theta$ 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 $f_2 + \epsilon(\tilde{f}_2 - f_2)$,

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

Hence,

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

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

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

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 $\frac{1}{b(\theta)}\varphi(\theta,\hat{\theta})=A'(m(\theta))\mathbf{1}_{u(\hat{\theta},\theta)\le 0}-\kappa(\theta,\hat{\theta})$ 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. □

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 $m(\theta) < m(\tilde{\theta})$. □

Proof of Claim 1. First, because $A : [0, 1] \to \mathbb{R}$ is concave, and is differentiable on $(0, 1]$ and globally continuous, it follows that $A'(0) \in [0, \infty]$ exists too, and the function $A' : [0, 1] \to \mathbb{R}_+ \cup \{\infty\}$ 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 $A'\left(G(\theta_{n-1}) - G(\theta_n)\right) = c - \theta_n$ 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 $A'\left(G(\theta_{n-1})\right) = \infty$—that is, if $G(\theta_{n-1}) = 0$ and $A'(0) = \infty$—then every $\theta \in (-\infty, \theta_{n-1}]$ has $A'\left(G(\theta_{n-1}) - G(\theta)\right) = \infty > c - \theta$, and so $\theta_n = -\infty$ is as desired. So let us further focus on the case that $A'\left(G(\theta_{n-1})\right) < \infty$.

In this case, define the continuous function

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

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, $c - (-\infty) < A'\left(G(\theta_{n-1})\right)$, so that this θ_n uniquely solves the equation, completing the inductive step.

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

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

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 $A'(0) \le c$. As noted above, θ has zero mass of barterers among types $\hat\theta > 0$. Moreover, any type $\hat\theta < 0$ has $c - \hat\theta > c \ge A'(0) \ge A'(m(\theta))$; hence, θ attention barters with zero mass of types $\hat\theta < 0$.

Now, we turn to the case that $A'(0) \ge c$. 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

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

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

$\mathbf{1}_{\theta>\tau(\hat{\theta}) \text{ and } \hat{\theta}>\tau(\theta)}$. 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 $\int f_2(\theta, \cdot)\, \mathrm{d}\mu = m^* = \int \mathbf{1}_{\check{\Theta}_n}\, \mathrm{d}\mu$, and hence $f_2(\theta, \cdot) =_{\text{a.e.}} \mathbf{1}_{\check{\Theta}_n}$. Morevoer, 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,

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

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

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

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

and

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

Hence, that $\theta \mapsto A'\left(G(\theta_0) - G(\underline{\theta})\right) - [c - \theta]$ 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

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

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

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

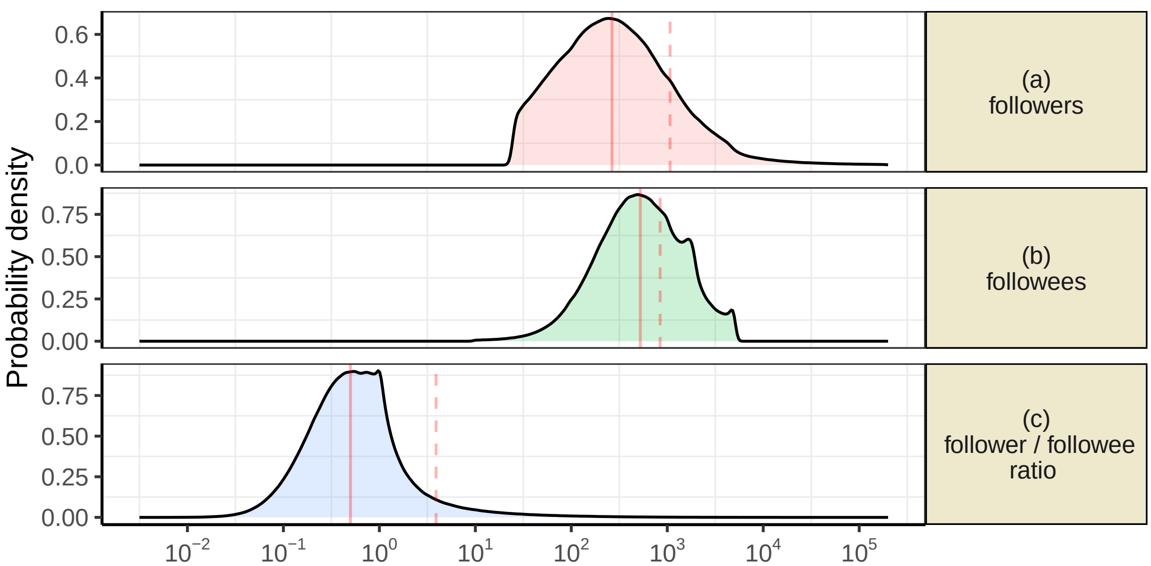

Notes: This figure reports kernel density estimates of the probability density function of the distributions of followers, followees, and follower-to-followee ratios in the entire data. For each facet, the vertical red lines depict the median (solid) and the mean (dashed) of the corresponding distribution, 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.

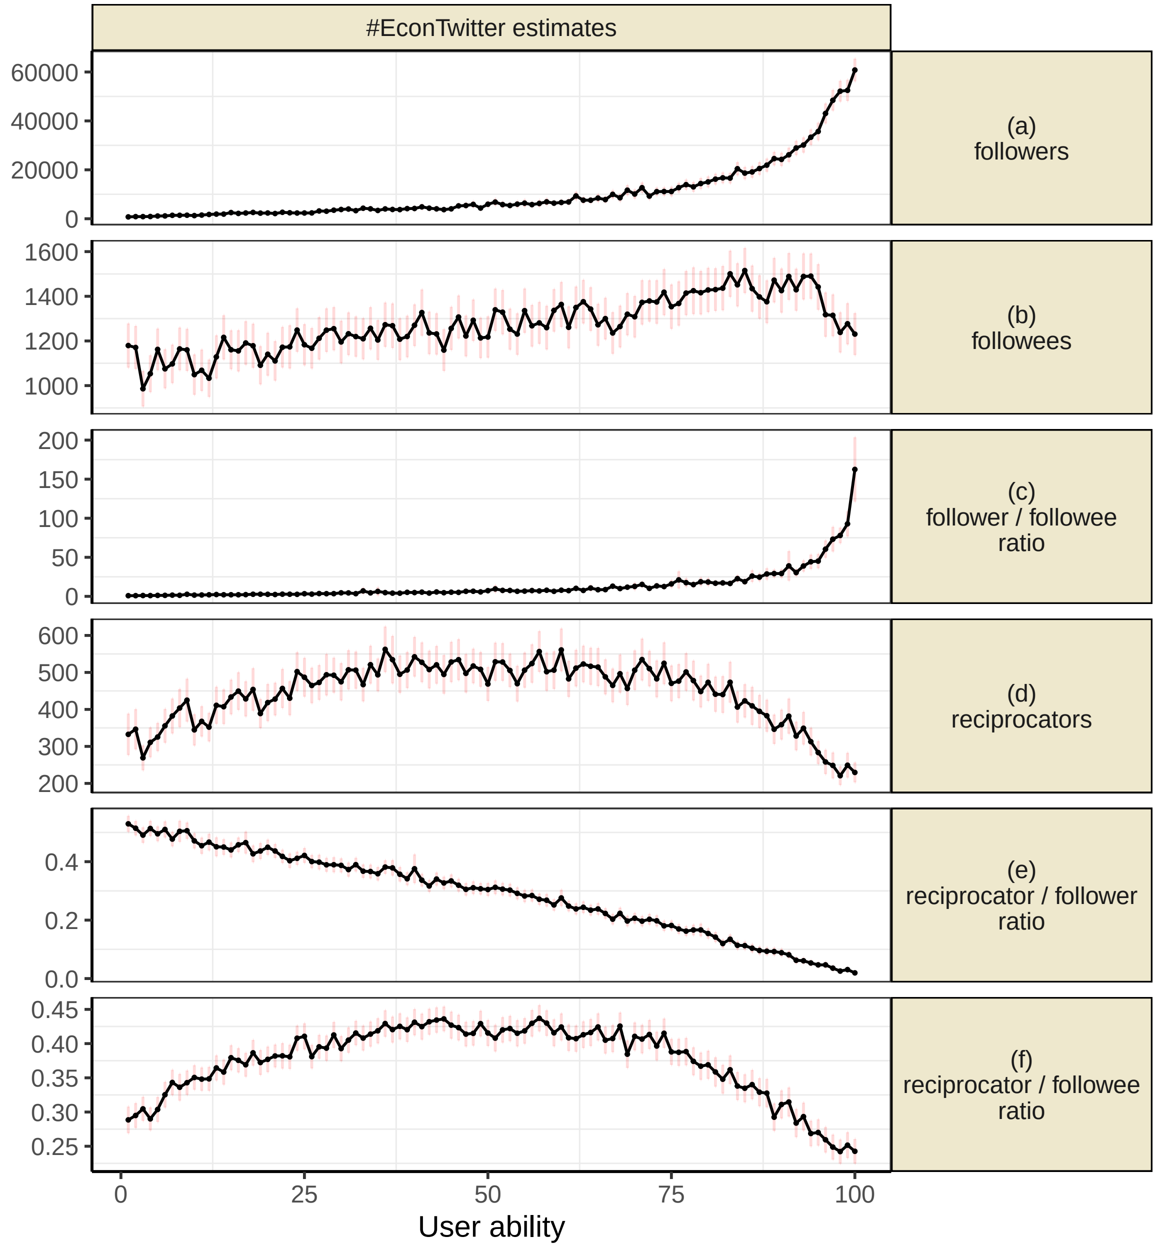

Notes: This figure reports unscaled #EconTwitter estimates of various network statistics, as a function of user ability. For more details on the construction of the estimates, see the description of Figure 4.

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.

	Dependent variable:		
	Ability (percentile)	Follower/followee Ratio	
	(1)	(2)	(3)
Tenure (days)	0.006***	0.0003***	
	(0.0001)	(0.00001)	
Ability (percentile)			0.025***
			(0.0001)
Constant	33.463***	0.777***	0.267***
	(0.375)	(0.016)	(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 (columne 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: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, 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

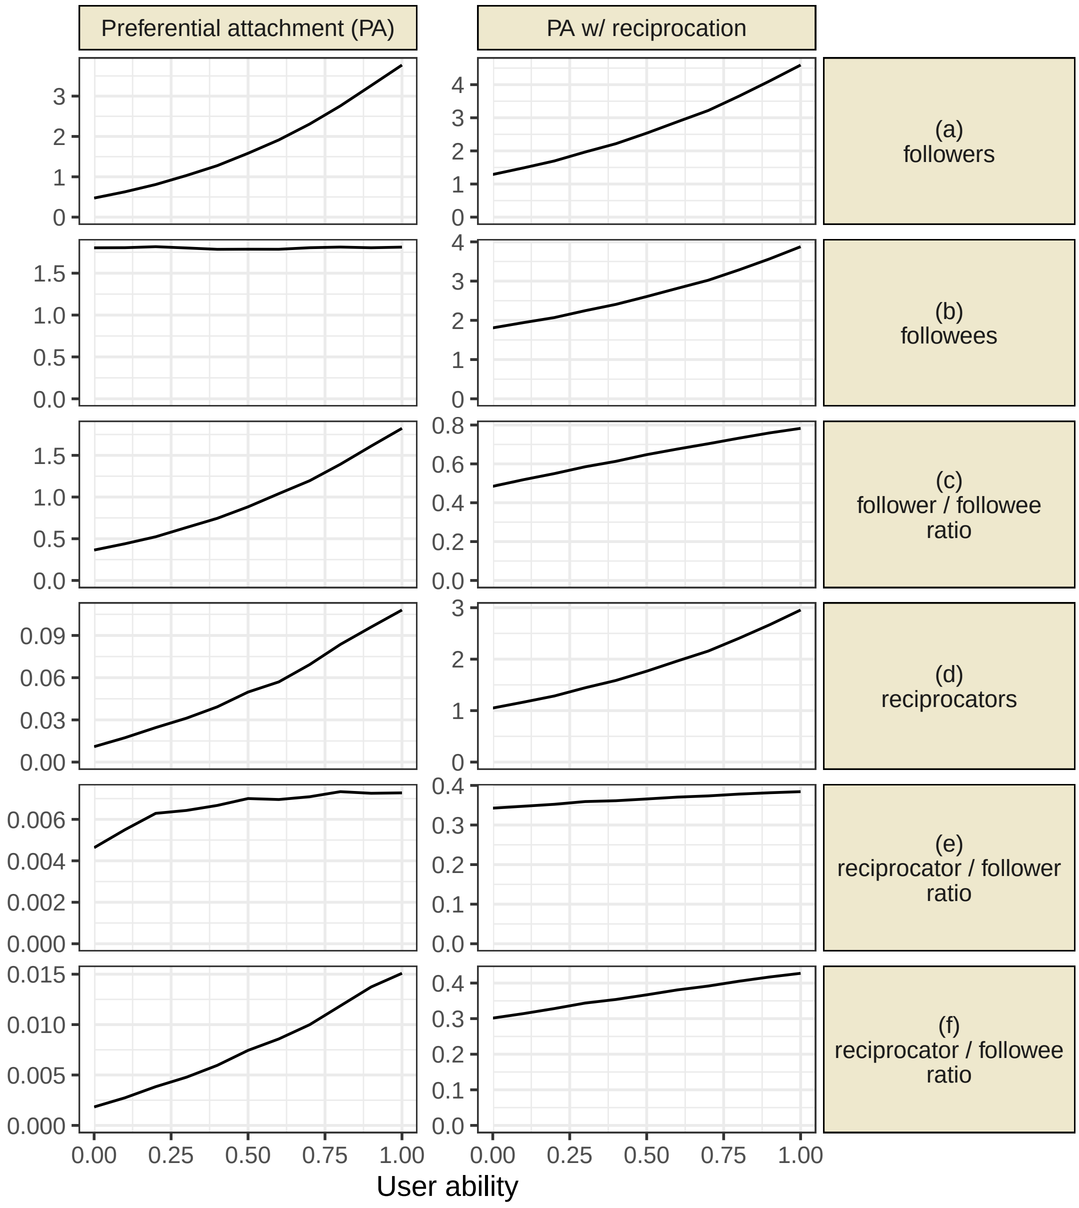

Notes: This figure reports the predictions of networks statistics using alternate network formation models. The left panel plots the predictions of the preferential attachment model (Bollobás et al., 2003) with parameters: $\alpha = 0.3$, $\beta = 0.5$, $\gamma = 0.2$, $\delta_{in} = 1$, $\delta_{out} = 1$. 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 Profific. 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
To what extent do you agree/disagree with the following statement:					
“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
When you check someone’s Twitter profile page, what do you typically pay attention to? (select all that apply)	
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
	More than 5 years	0.711
How often do you use Instagram?	Once a month or less	0.031
	A few times a week	0.122
	Once a day	0.132
	Multiple times a day	0.715
How many Instagram followers do you have?	Less than 50	0.168
	50 to 500	0.515
	501 to 2,000	0.253
	2,001 to 10,000	0.059
	More than 10,000	0.004
How many users do you follow on Instagram?	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
To what extent do you agree/disagree with the following statement:					
“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