MANAGEMENT SCIENCE
Articles in Advance, pp. 1–16
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)

http://dx.doi.org/10.1287/mnsc.2016.2476

Customer Referral Incentives and Social Media

Ilan Lobel
Stern School of Business, New York University, New York, New York 10012, ilobel@stern.nyu.edu

Evan Sadler
Harvard University, Cambridge, Massachusetts 02138, esadler@fas.harvard.edu

Lav R. Varshney
Department of Electrical and Computer Engineering, University of Illinois at Urbana–Champaign, Urbana, Illinois 61801, varshney@illinois.edu

We study how to optimally attract new customers using a referral program. Whenever a consumer makes a purchase, the firm gives her a link to share with friends, and every purchase coming through that link generates a referral payment. The firm chooses the referral payment function and consumers play an equilibrium in response. The optimal payment function is nonlinear and not necessarily monotonic in the number of successful referrals. If we approximate the optimal policy using a linear payment function, the approximation loss scales with the square root of the average consumer degree. Using a threshold payment, the approximation loss scales proportionally to the average consumer degree. Combining the two, using a linear payment function with a threshold bonus, we can achieve a constant bound on the approximation loss.

Keywords: social networks; pricing in networks; microeconomic behavior; stochastic networks
History: Received March 24, 2015; accepted January 29, 2016, by Noah Gans, stochastic models and simulation. Published online in *Articles in Advance* July 14, 2016.

1. Introduction

Referrals have emerged as a primary way through which companies in the social networking era acquire new customers. Instead of spending money on traditional advertising, many companies now rely on social-media-based referral programs to bring in new customers. Referral programs are used throughout the economy, but their use is particularly common among start-ups, where they are seen as both an affordable and an effective way to grow.

Referral programs come in many forms. Some companies pay customers for every new referral. Others pay customers only if they bring in sufficiently many referrals. We propose a framework to analyze different referral program designs and determine good payment rules. There are two main parts to our analysis. First, we identify the optimal payment function; for each possible number of successful referrals that a consumer makes, we find the optimal reward the firm should pay. Second, we examine simple and easy-to-implement payment functions—linear, threshold, and a combination of the two—to assess how well they can approximate the profits of the optimal policy.

Two companies that exemplify the success of referrals as a customer acquisition strategy are Living Social and Dropbox. Living Social is a daily deal website that offers discounted prices for a variety of goods and services. Living Social was launched in 2007 and achieved a multibillion dollar valuation only four years later. At the heart of Living Social's rapid growth was its viral marketing strategy. Whenever someone buys a deal from them, Living Social encourages the customer to post the transaction on social media. If at least three friends subsequently purchase based on this post, Living Social provides the customer with a full refund of the purchase price.

Dropbox, a cloud-based file hosting service that was founded in 2008, also achieved extraordinarily fast growth through referrals. Dropbox offers every user 2 GB of free storage, and it increases this amount by 500 MB for every friend the user refers (or 1 GB if a friend signs up for a professional account), up to a maximum of 16 GB for free. The CEO of Dropbox attributes much of the company's early success, going from a hundred thousand to four million users in just 15 months, to its referral program (Houston 2010). Many other technology companies have grown through referral programs, including the electronic payments company PayPal in the early 2000s, the Living Social competitor Groupon, and the car service companies Uber and Lyft.

What is the best way to structure referral payments? Linear payments, whereby every successful referral generates the same payment, are the simplest and probably the most common model. Dropbox uses linear payments capped at 28 referrals, and Uber offers twenty dollars per successfully referred friend. In contrast, Living Social uses a threshold model in which two referrals yield no benefit, but three lead to a full refund. In principle, the optimal referral payments could follow a linear function, a threshold function, or any other nonlinear payment function.

We consider a firm that designs its referral incentive program with two objectives: extract immediate revenue and advertise to potential customers. We focus on the interaction between this firm and a focal consumer currently deciding whether to make a purchase. If the consumer buys the product, she can refer as many of her friends as she wants. The firm wishes to maximize a linear combination of the revenue obtained from the focal consumer and the value obtained from her referrals. The consumer plays an equilibrium of a network game in which she balances the costs and rewards associated with making referrals, taking account of the referral program's value to her friends. When a referred friend purchases the firm's product, we say that a conversion occurs.

An important assumption is that the firm values referrals, both those that convert and those that do not, but only pays consumers for conversions. We make this assumption because the advertising value of a referral program goes beyond immediate sales, including product and brand awareness, as well as future sales. Even though the firm values referrals, we assume that the firm chooses to pay consumers only for conversions. One reason is that measuring referrals is more difficult than measuring purchases. More significantly, paying for referrals directly creates opportunities to cheat. Customers could send referrals to fake email accounts, earning rewards without creating any value for the firm. Requiring a purchase for the referral to count helps to avoid this problem. We consider the case in which conversions and nonconversions are valued differently in Section 6.1.

Our first main result characterizes the optimal referral policy when there are no restrictions on the payment function the firm can use. A key part of our analysis is a decomposition of the expected referral payment into two parts: one that compensates consumers for the cost of making referrals and the remainder, which is the true reward or discount that the consumer receives. Since high-degree consumers can bring in more referrals, they optimally receive larger rewards, making them more likely to purchase. Although the optimal expected reward is increasing in the consumer's degree, the optimal payment function might be nonmonotonic in the number of conversions. This nonmonotonicity occurs because the firm values referrals but pays for conversions. The number of conversions is a stochastic function of the number of referrals, so the expected reward a consumer receives is a smoothed version of the payment function, and precise control of the expected reward may require a complex payment rule. This implies that the optimal policy is, at least in some instances, quite complicated and potentially impractical to implement. Few if any firms would ever offer a deal like "bring three friends and earn \$20, or bring four friends and earn \$15" as consumers would likely find such programs confusing and would have difficulty responding optimally.

The nonmonotonicity of optimal payments motivates our study of simpler, more realistic payment functions. We measure performance by benchmarking these payment functions against the returns from using optimal payments. Linear payment functions provide an obvious starting point. To evaluate how good of an approximation we can achieve with a linear policy, we perform a worst-case analysis with respect to the degree distribution of the network. We show that worst-case losses from using the best linear incentive functions scale according to the square root of the average degree of the social network. Another natural approximation is a threshold payment function, though these generally fare much worse than linear policies. With a threshold payment function, the worst-case losses relative to the optimum scale linearly with the average degree in the network.

The different scaling rates for linear and threshold policies give an incomplete comparison because losses in each case come from very different sources. When the degree distribution has a heavy tail, linear policies are prone to overpaying high-degree consumers. Threshold policies provide a cap on payments, but they have difficulty properly compensating consumers for the costs of making referrals. To render these insights more concretely, we present a numerical analysis of a few special cases. Comparing the performance of a linear payment function across Poisson, geometric, and power-law degree distributions, we can clearly see how tail thickness negatively impacts profits. In social-media settings, in which degree distributions typically follow a power law and the marginal cost of making a referral is low, threshold payments perform relatively well. In fact, if the marginal cost of making a referral is zero, we can show that the losses from using a threshold payment are bounded by a constant across all possible degree distributions.

Since these two types of policies have complementary strengths, combining them may offer a significant improvement over either one alone. Making linear payments, with an added threshold bonus, captures the best features of both policies. The linear part compensates consumers for the social costs of making referrals, and the threshold bonus can provide the desired discount for high-degree consumers without overpaying. This combination attains a constant bound on the approximation loss, regardless of the degree distribution or the cost of making referrals.

Our results offer a useful approach to referral program design when there is limited information about consumer preferences and the social network structure. Our main results assume that the firm knows only the valuation and degree distributions and that consumer knowledge, beyond a consumer's own type, is similarly limited. In practice, firms and consumers may have additional information. Our broader contribution is an adaptable framework on which to build extensions. In Section 6, we consider several potential directions to explore, including the problem of designing incentives when consumers know whether their neighbors have low or high degrees. We show that even though the optimal design problem is harder to solve to optimality, under some conditions we can guarantee higher firm profits relative to the case in which consumers have no information about their neighbors' degrees.

1.1. Related Literature

Personal referrals have a powerful effect on consumers. Marketing researchers, using the term "word of mouth," have studied the impact of referrals for decades. This work credits interpersonal communication with far greater influence over consumer attitudes and behavior than either conventional advertising or neutral print sources (Buttle 1998), and the value of personal referrals can constitute a significant portion of a customer's value (Kumar et al. 2007). With the growth of online platforms, and social media in particular, we can now measure these effects more precisely than ever before. For instance, Godes and Mayzlin (2004) are able to measure word-of-mouth effects using conversations from an online discussion forum. Studies in both the computer science literature and the information systems literature draw on data from online social networks to assess just how much people influence one another in different settings (Leskovec et al. 2007, Aral and Walker 2011). There is also work in experimental economics showing that social learning effects on demand are at least as big as traditional advertising channels (see Mobius et al. 2011).

Parallel work asks questions about design. Algorithmic and simulation-based approaches allow researchers to study the relative importance of weak and strong ties (Goldenberg et al. 2001), optimal seeding strategies (Kempe et al. 2003), and how to use viral product features to maximize the spread of adoption in a network (Aral et al. 2013). Many papers in the economics literature explore word-of-mouth communication as a signaling game in which consumers learn about product quality through friends, and the firm manipulates the information consumers receive, either through hired "trendsetters" (Chatterjee and Dutta 2014) or launch strategies (Campbell 2015). Campbell (2013) studies pricing in the presence of word-of-mouth learning. Our firm more directly manipulates word-of-mouth information transmission through its referral program.

Our equilibrium analysis draws on the recent economics literature on games with local network externalities (Ballester et al. 2006, Sundararajan 2007, Galeotti et al. 2010). A referral program provides a way to artificially create network externalities, inducing complementarities in friends' purchase decisions; an equilibrium of our consumer game is mathematically equivalent to an equilibrium of a network game with strategic complements. Candogan et al. (2012), Bloch and Quérou (2013), and Cohen and Harsha (2015) build on these results to create a theory of optimal pricing for a monopolist selling goods to networks of consumers. Our paper brings a new perspective to this growing literature, taking a mechanism design approach that provides insight on how to structure the network externality in order to generate the desired information transmission.

A few other papers take up the problem of referral payment design. Biyalogorsky et al. (2001) consider how the firm can jointly optimize pricing and referral incentives over the lifetime of a consumer. We introduce two important innovations on this approach. First, our consumers strategically anticipate the value that their friends will derive from the referral program. Second, we address the role of the social network structure. Libai et al. (2003) study the somewhat different setting of affiliate marketing, allowing a richer information set and individualized referral contracts. As typical customer referral programs cannot implement individualized referral contracts or keep track of leads separately from conversions, we restrict our attention to pay-per-conversion systems.

Our approach necessarily abstracts away from many details of referral programs that are nevertheless important. The medium through which customers communicate with each other has implications both for the cost of making referrals and the value the firm derives from referrals. Chu and Kim (2011) note that the influence people exert through online interactions is different than that through offline interactions, and Burke and Kraut (2014) present evidence from Facebook suggesting that one-on-one communication is fundamentally different than multicast communication. Although we do not explicitly address what medium is best for a referral program, we provide a framework to think about this question. The type of interactions a firm relies upon to generate referrals will determine the cost, payoff, and network parameters that are primitives in our model, allowing a comparison between optimal policies using different modes of communication.

2. The Referral Game

A firm produces a product at zero marginal cost and sells it at a fixed price $p > 0$.[1] We consider a countable population of potential consumers, which we represent as a rooted graph. We refer to the root of this graph as the focal consumer. The focal consumer is deciding whether to purchase at the current moment. The firm uses a referral incentive program to expand its customer base and increase profits. When the focal consumer makes a purchase, the firm offers her a code or a link to share with friends. When a friend uses that code or link to make a purchase, the original consumer may receive a reward from the firm.

There are two distinct parts to our model. First, the firm offers a referral program to the focal consumer and derives a payoff based on whether the focal consumer purchases and how many referrals get sent. The second part is a game played among consumers: the focal consumer must predict how many of her friends will respond to a referral, and to do this she must account for the referral program's value to her friends.

A referral program is a payment function $\mathbf{w}: \mathbb{N} \to \mathbb{R}$ specifying a monetary payoff for each possible number of successfully referred friends. Aside from the restriction $\mathbf{w}(0) = 0$, we allow the firm to choose any payment function. The payment is a function of the number of friends who purchase in response to a referral, not the number of referrals sent. We write $\mathcal{W}$ for the space of possible referral payment functions. After observing the firm's referral program, the focal consumer decides whether to buy, and if so, how many friends to refer. The focal consumer has some valuation $v \in [0, 1]$ for the product and has d friends she could potentially refer. She chooses a pair $\mathbf{s} = (b, r)$, where $b \in \{0, 1\}$ indicates whether the consumer purchases the product ($b = 0$ represents no purchase), and $r \in \{0, 1, 2, \ldots, d\}$ represents the number of friends who are referred. The focal consumer incurs a cost $\delta \geq 0$ for making each referral, and the firm earns $y > \delta$ for each referral that gets sent. We interpret y as the expected value of a referral made to a random consumer. The value $y - \delta > 0$ is the expected surplus generated from a referral.

To decide whether to buy and how many referrals to make, the focal consumer must form expectations about how her friends will respond to a referral. Each consumer has a valuation $v \in [0, 1]$ for the product drawn independently from a continuous distribution F, which has density f. The focal consumer has degree d drawn according to the distribution G, and we use g_d to denote the probability of having d friends. The friends of the focal consumer have degrees drawn independently according to the size-biased distribution G' with

$$\mathbb{P}(G' = d) = \frac{d\,\mathbb{P}(G = d)}{\mathbb{E}[G]}, \tag{1}$$

but if we suppose that these friends can only refer new potential consumers, the effective degree d of a friend is distributed according to G'' with $\mathbb{P}(G'' = d) = \mathbb{P}(G' = d + 1)$.[2] The distributions F, G, and G'' are common knowledge. A consumer's information set comprises her valuation and degree pair (v, d); consumers do not observe the valuations or the degrees of their neighbors. A friend of the focal consumer (and a friend of a friend, and so on) faces the same decision problem as the focal consumer—that is, the same payment function $\mathbf{w}$, the same referral cost δ, and the same beliefs about friends' valuations and degrees.

Each consumer who becomes aware of the product through a referral chooses a pair (b, r) just as the focal consumer does; the strategy space for a consumer with d friends is $A_d = \{(0, 0)\} \cup \{(1, r) \mid r = 0, 1, \ldots, d\}$. A pure strategy for a consumer is a function $\sigma(v, d, \mathbf{w})$: $[0, 1] \times \mathbb{N} \times \mathcal{W} \to A_d$ specifying an action for each possible valuation, degree, and payment function. We use b_σ and r_σ, respectively, to denote the two components of σ, and we use Σ to denote the space of pure strategies.

The firm's payoff depends on the actions of the focal consumer.[3] Let the random variable $N_\sigma(r)$ denote the number of neighbors who respond to a referral when the focal consumer refers r neighbors and other consumers use the strategy σ. If the focal consumer chooses (b, r), the firm earns

$$\pi_{(b,\, r)} = b(p - \mathbb{E}[\mathbf{w}(N_\sigma(r))] + ry).$$

If the consumer buys the product, the firm earns the purchase price p, less the expected referral payment, plus an additional payoff that is proportional to the number of referrals.

We consider equilibria of the consumer game in pure symmetric strategies, using σ' to denote a unilateral deviation when all other consumers play σ. The strategy σ induces expectations over the distribution of $N_\sigma(r)$. Playing σ' in response to others playing σ yields an expected payoff

$$u_{\sigma',\sigma}(v, d, \mathbf{w}) = b_{\sigma'}(v - p - \delta r_{\sigma'} + \mathbb{E}_\sigma[\mathbf{w}(N_\sigma(r_{\sigma'}))]). \tag{2}$$

[1] For now, we take the price as exogenous, but we address the issue of optimal pricing in Section 6.2.

[2] This size-biased distribution is a standard correction that generates the degree distribution conditional on being a person's neighbor (see, for instance, Jackson and Yariv 2007). This implicitly assumes that players' degrees are uncorrelated. We slightly abuse notation in Equation (1) by using G to also represent a random variable with distribution G.

[3] Alternatively, we could view this payoff in terms of the behavior of an average or representative consumer who becomes aware of the product.

A purchasing consumer earns her valuation for the product v, less the purchase price and the cost of making referrals, plus the expected referral payment. A strategy profile[4] $\sigma \in \Sigma$ is a pure strategy Bayesian equilibrium of the consumer game if

$$u_{\sigma,\sigma}(v,d,\mathbf{w}) \geq u_{\sigma',\sigma}(v,d,\mathbf{w}) \quad \text{for all } (v,d,\mathbf{w}) \in [0,1] \times \mathbb{N} \times \mathscr{W} \text{ and all } \sigma' \in \Sigma. \quad (3)$$

The firm first chooses a payment function $\mathbf{w}$. The focal consumer then observes this choice and plays a symmetric equilibrium for the consumer game in response. Our goal is to understand optimal or near-optimal payment functions in terms of the referral surplus value y, the consumer's cost of a referral δ, and the distributions F and G of consumer valuations and degrees. Formally, we study the firm's optimization problem:

$$\begin{aligned} \sup_{\mathbf{w}\in\mathscr{W},\, \sigma\in\Sigma} \quad & \mathbb{E}_{v,d}[b_\sigma(p - \mathbf{w}(N_\sigma(r_\sigma)) + r_\sigma y)] \\ \text{s.t.} \quad & u_{\sigma,\sigma}(v,d,\mathbf{w}) \geq u_{\sigma',\sigma}(v,d,\mathbf{w}) \\ & \text{for all } \sigma' \in \Sigma,\ v \in [0,1],\ d \in \mathbb{N}. \end{aligned} \quad (4)$$

We can interpret this model as a mechanism design problem, with the firm maximizing an objective subject to incentive compatibility and individual rationality constraints. We note that this mechanism design problem is two-dimensional, because every consumer has two pieces of private information: her own valuation and her degree.

2.1. Remarks on the Model

Our modeling choices reflect a difference in perspective between the consumers and the firm. For a consumer deciding whether to buy, the decision problem is fundamentally static: she is making a single decision in a given moment. The potential for future purchases or future interactions with the firm do not enter the consumer's thought process. Our focal consumer assumes that any friend she refers will face the same basic decision problem, including the referral program, so she accounts for the value of the referral program to her neighbors when she evaluates the value of the program to her. Even though friends make purchase decisions in the future, formally our consumers play a symmetric Bayes-Nash equilibrium of a one-shot simultaneous move game.

In contrast to the consumer, the firm takes a broader view of the problem: the firm is interested in growing its business. Beyond immediate purchases and referrals, the firm is interested in future referrals to friends of friends, repeat purchases, the possibility of changing the referral program over time, and building brand awareness and loyalty. To abstract away from this complexity, we focus on a single interaction between the firm and a focal consumer, and we incorporate all of these elements into the expected referral value y. We think of y as the expected value of all future consequences of a referral. In principle, this could emerge as a continuation value in a dynamic game in which the firm updates its referral program as the customer base grows. However, adding this additional layer to the model would not change our structural results for optimal and approximately optimal payment functions at a particular instant. Alternatively, we can view the referral program as a substitute for traditional advertising. Under this interpretation, the referral value y is an outside option representing the cost of achieving equivalent reach through standard advertising channels. Taking either view, we can understand important trade-offs that firms face in referral program design.

We summarize the network structure through the degree distribution G. The consumers effectively treat the network as an infinite tree in which every neighbor has a distinct and independent collection of neighbors in turn. To generate such beliefs, we could take the large network limit of a configuration model: for a fixed population size N, generate the network uniformly at random from all possible networks with a given degree sequence, and select the root node uniformly at random. As N approaches infinity, beliefs in such a model converge to those of our consumers.

Although real networks do not follow this structure, our representation admits at least two interpretations. One focuses on consumers' cognitive limitations: they adopt the infinite tree model as a subjective representation of the network because it simplifies decision making. Taking full account of the network before making a purchase presents a very complex decision problem. Our assumptions about consumer behavior retain the fundamental trade-offs in the consumer problem while reducing the computational burden. The "infinite tree model" can also present a decent approximation to actual diffusion paths in a large network. These paths are trees, trees that emerge if we prune cycles and ignore redundant referrals. In a large network, the proportion of people who have already purchased changes little moving from a consumer to her immediate neighbors, so the choice problems should be quite similar.

Many of the assumptions that simplify our exposition are not crucial to the main results. People typically find out about a product through multiple channels, and they may not bother to read every referral offer they receive. In the context of our model, we could assume that some friends who receive referrals have already purchased the product, or they may simply ignore the

[4] Note the slight abuse of notation: we use the same representation Σ to represent symmetric pure strategy profiles that we use to represent a consumer's space of pure strategies.

offer, with some independent probability; this would have no impact on our qualitative findings. Looking at the expression for consumer utility in Equation (2), this change would appear in the random variable $N_\sigma(r_{\sigma'})$. This random variable follows a binomial distribution with $r_{\sigma'}$ trials and a success probability that is a function of the strategy profile σ. We can discount this success probability by some factor $\alpha < 1$ without changing the form of either the consumer optimization (3) or the firm optimization (4).

If we wish to incorporate clustering, reflecting that friends tend to have friends in common, we could add a dampening factor to the model and assume that our neighbors' neighbors have a proportionally lower chance of responding than our neighbors do. The problems (3) and (4) again would not change, but we would expand the consumer's information set to include the probability α that a neighbor is at all receptive to a referral. A strategy profile σ would become a function of the variables $(v, d, \mathbf{w}, \alpha)$, and a consumer who believes neighbors are receptive with probability α would believe that neighbors of neighbors are receptive at some lower probability $\beta\alpha$, with $0 < \beta < 1$. Though imperfect, adjustments like these can give a reasonable first-order approximation for what happens in more realistic networks.

We can also relax the linear structure of referral benefits and costs. From the firm's perspective, it would make sense to model declining marginal returns from referrals, especially with clustering in the consumer network. Versions of all our results continue to hold under mild conditions.[5] Our results are robust to adding a fixed cost for making any referrals on top of the cost δ per referral.

3. Consumer Behavior

We first study the consumer equilibria for a fixed referral payment function $\mathbf{w} \in \mathcal{W}$. Throughout this section, we take $\mathbf{w}$ as given, and we suppress dependence on $\mathbf{w}$ in our expressions to simplify notation. Suppose the focal consumer believes that each neighbor will purchase in response to a referral independently with some probability P. The true likelihood that a neighbor will purchase is determined in equilibrium—it is an expectation over behavior as a function of valuations and degrees—but for now take this P as given. This neighbor purchase probability, together with the number of neighbors the consumer has, determines the referral incentive's value for the consumer. Let $B_P(r)$ denote a binomial random variable with r trials and success probability P. If a consumer has d neighbors, each of whom purchases in response to a referral with probability P, the value of the referral program is

$$I_P(d) = \max_{r \in \{0,1,\ldots,d\}} \mathbb{E}[\mathbf{w}(B_P(r)) - r\delta]. \tag{5}$$

Combining Equations (2) and (5), we see that a focal consumer with d neighbors will purchase if and only if her valuation v is at least $p - I_P(d)$. The threshold

$$v_P(d) = p - I_P(d) \tag{6}$$

is the lowest valuation at which a consumer with degree d will buy, given that a neighbor will purchase with probability P. If other players' strategies result in a neighbor purchase probability P, then the thresholds $v_P(d)$ define a best reply strategy. This best reply depends on the other players' strategies solely through the neighbor purchase probability P.

A random neighbor with degree d who best responds to the purchase probability P will in turn purchase with probability $1 - F(v_P(d))$. Taking an expectation over all possible degrees gives the neighbor purchase probability that results from neighbors playing a best reply to P:

$$\phi(P) = \sum_{d \in \mathbb{N}} g''_d (1 - F(v_P(d))). \tag{7}$$

The function ϕ allows us to give an alternative characterization of equilibrium strategy profiles. Instead of working directly with utility functions as in Equation (3), we can work with neighbor purchase probabilities. There is a correspondence between equilibrium strategy profiles and fixed points of the map ϕ.

Given an equilibrium strategy profile σ, neighbors purchase independently with some probability P_σ, and we must have $\phi(P_\sigma) = P_\sigma$. If P_σ is not a fixed point, then a random consumer best responding to P_σ would purchase with a different probability, which is inconsistent with P_σ being an equilibrium. Conversely, for any P^* that is a fixed point of ϕ, the thresholds $v_{P^*}(d)$ define an equilibrium strategy uniquely up to a measure zero set of valuations who are indifferent about purchasing. The correspondence between equilibria and fixed points of ϕ yields a simple existence proof.

Proposition 1. *A symmetric pure strategy Bayesian equilibrium of the consumer game exists.*

Proof. This amounts to showing that a fixed point exists for the function ϕ defined in Equation (7). Since $I_P(d)$ is a continuous function of P, and F is a continuous distribution, the map ϕ is continuous. The existence of a fixed point is an immediate consequence of the intermediate value theorem. □

[5] Without further assumptions, we can still decouple the firm's optimization problem as in Theorem 2 and achieve perfect price discrimination across consumer degrees. The nonmonotonicity result in Proposition 3 will hold as long as some consumers with many friends should receive the product at an effective price of zero under the optimal policy. The asymptotic results of Section 5 continue to hold as long as the marginal value of a referral is bounded above δ.

This result establishes existence but not uniqueness of the equilibrium. Some parameter values and payment functions support multiple neighbor purchase probabilities in equilibrium. For instance, suppose that $p > 1$, that $g_0'' = 0$, and that $\mathbf{w}(d) > p$ for each $d > 0$. In this case, both $P_\sigma = 0$ and $P_\sigma = 1$ are equilibrium purchase probabilities. The former is an equilibrium in which no consumer purchases because the good is too expensive and the referral payment contributes no value since others never respond to referrals. The latter resembles a pyramid scheme in which all consumers buy the good only because they expect referral payments to entirely cover the cost.

However, as long as the referral payments are not too generous, pyramid scheme equilibria will not arise, and we should expect a unique equilibrium. A sufficient condition for uniqueness is that $\phi(P) - P$ is a decreasing function of P. Equivalently, the equilibrium is unique if $|\phi(P_1) - \phi(P_2)| < |P_1 - P_2|$ for all $P_1, P_2 \in [0, 1]$. The following result provides a sufficient condition for uniqueness.

Proposition 2. *Let* $\bar{w} = \sup_d(\mathbf{w}(d) - \mathbf{w}(d-1))$, *and let* $\bar{f} = \sup_{v\in[0,1]} f(v)$. *If we have* $\bar{f} \cdot \bar{w}\,\mathbb{E}_{D\sim G}[D] < 1$, *then the pure strategy Bayesian equilibrium of the consumer game is unique.*

Proof. We compute

$$\begin{aligned}
|\phi(P_1) - \phi(P_2)| &\le \sum_{d\in\mathbb{N}} g_d |F(v_{P_2}(d)) - F(v_{P_1}(d))| \\
&\le \bar{f} \sum_{d\in\mathbb{N}} g_d |v_{P_2}(d) - v_{P_1}(d)| \\
&= \bar{f} \sum_{d\in\mathbb{N}} g_d |I_{P_1}(d) - I_{P_2}(d)| \\
&\le \bar{f} \sum_{d\in\mathbb{N}} g_d d\bar{w} |P_1 - P_2| \\
&= \bar{f} \cdot \bar{w}\,\mathbb{E}_{D\sim G}[D] |P_1 - P_2|.
\end{aligned}$$

The last inequality is the critical step. A consumer with d neighbors can gain no more than $\bar{w}|P_1 - P_2|$ per neighbor when the neighbor purchase probability changes. The conclusion follows. □

The condition in Proposition 2 gives us a bound $\bar{w}$ on the incremental benefit from making one more successful referral. To get an intuitive sense of how to apply the condition, suppose F is uniform so that $\bar{f} = 1$. In this case, the condition is satisfied if a consumer with an average number of neighbors can never obtain a total referral payment of 1, the highest amount that any consumer would be willing to pay for the product. Of course, the uniform distribution is a best-case scenario with respect to all distributions with support in $[0, 1]$.

4. Optimal Incentives

We now turn to the firm's problem of selecting the optimal payment function $\mathbf{w}$, and we reintroduce the dependence on $\mathbf{w}$ in our expressions. We can decompose the firm's referral payments into two components: part of the payment compensates consumers for making referrals, and the rest provides a reward (or an indirect discount) on the purchase price. If a consumer with d neighbors makes $r_\sigma(d)$ referrals in equilibrium σ, we can rewrite Equation (5) as

$$\mathbb{E}_\sigma[\mathbf{w}(B_{P_{\sigma,\mathbf{w}}}(r_\sigma(d)))] = I_{P_{\sigma,\mathbf{w}}}(d) + r_{\sigma,\mathbf{w}}(d)\delta. \tag{8}$$

The value of the incentive $I_{P_{\sigma,\mathbf{w}}}(d) \ge 0$ to a consumer is equivalent to a reduction in the effective price, and $r_{\sigma,\mathbf{w}}(d)\delta$ is the amount the firm must pay in expectation to compensate the consumer for making referrals.

There are at least two reasons to use a referral incentive to discount the product beyond compensating for referrals. First, the referral incentive offers a way to price discriminate between consumers of different degrees. Since the surplus from referrals $y - \delta$ is positive, consumers with more neighbors can generate more surplus if they make purchases. Intuitively, an optimal incentive should offer these consumers a lower effective price so that more of them buy. Second, if the good is subject to a Veblen effect, a referral incentive could provide a way to retain the demand associated with a high list price while discounting the good to an optimal effective price.

Decomposing referral payments in this way allows for a simpler representation of the firm's decision problem. The firm can fully describe consumer behavior using a sequence of referral and discount pairs $\{(r_d, I_d)\}_{d\in\mathbb{N}}$. For a consumer with d neighbors, the value r_d specifies how many referrals she makes conditional on a purchase, and I_d is the discount in equilibrium she obtains, which determines how likely the consumer is to make a purchase. The firm's problem is then to choose a sequence $\{(r_d, I_d)\}$ that will maximize profits. The firm is constrained to choose a sequence it can implement in some equilibrium of the consumer game. We can restate the firm's problem (4) as

$$\begin{aligned}
&\max_{\{(r_d, I_d)\}} \sum_{d\in\mathbb{N}} g_d(1 - F(p - I_d))(p - I_d + r_d(y - \delta)) \\
&\text{s.t.}\quad \{(r_d, I_d)\} \quad \text{implementable.}
\end{aligned} \tag{9}$$

Implementability necessitates that the sequence of referral and discount pairs satisfies individual rationality and incentive compatibility constraints. Feasibility requires that $r_d \in \{0, 1, \ldots, d\}$ and individual rationality mandates that $I_d \ge 0$ for each d. Since a consumer with many neighbors can always mimic one with fewer neighbors, incentive compatibility implies the following:

(a) The r_d are nondecreasing in d, and whenever $r_d \ne r_{d-1}$ we have $r_d = d$.

(b) The I_d are nondecreasing in d, and $I_d > 0$ only if $r_d > 0$.

(c) Whenever $r_{d'} = r_d$, we have $I_{d'} = I_d$.

It turns out that satisfying the individual rationality and incentive compatibility constraints is sufficient to ensure that a sequence is implementable: for any sequence satisfying the above properties, there is a corresponding payment function $\mathbf{w}$ and an equilibrium σ of the consumer game producing the given sequence.

Theorem 1. *A sequence of referral and discount pairs* $\{(r_d, I_d)\}_{d\in\mathbb{N}}$ *is implementable if and only if the individual rationality and incentive compatibility constraints* (a), (b), *and* (c) *are satisfied.*

Proof. Given a sequence of referral and discount pairs, an equilibrium neighbor purchase probability is uniquely determined as

$$P_{\sigma,\mathbf{w}} = \sum_{d\in\mathbb{N}} g''_d(1 - F(p - I_d)).$$

Assuming this neighbor purchase probability, we can inductively construct a payment function $\mathbf{w}$ that implements the corresponding referral and discount pairs. Note that the discount to customers of degree d can only depend on the values of $\mathbf{w}(k)$ for $k \le d$.

For the case $d = 1$, if $r_1 = 0$ take $\mathbf{w}(1) = 0$, otherwise set $\mathbf{w}(1) = (I_1 + \delta)/P_{\sigma,\mathbf{w}}$. Suppose we have constructed a $\mathbf{w}$ that implements a given sequence up to customers of degree d, and consider two cases. First, if $r_{d+1} \ne d+1$, we have $r_{d+1} = r_d$ and $I_{d+1} = I_d$. Setting $\mathbf{w}(d+1)$ sufficiently low will make the value of referring the $(d+1)$th neighbor negative, so these consumers will mimic those of degree d as desired. Alternatively, we have $r_{d+1} = d+1$ and $I_{d+1} \ge I_d$. We wish to define $\mathbf{w}(d+1)$ so that

$$I_{d+1} + (d+1)\delta = \sum_{j=0}^{d+1} \binom{d+1}{j} (P_{\sigma,\mathbf{w}})^j (1 - P_{\sigma,\mathbf{w}})^{d+1-j} \mathbf{w}(j).$$

Since we have already defined $\mathbf{w}$ up through d, we can set

$$\mathbf{w}(d+1) = \frac{1}{(P_{\sigma,\mathbf{w}})^{d+1}} \cdot \left[I_{d+1} + (d+1)\delta - \sum_{j=0}^{d} \binom{d+1}{j} \cdot (P_{\sigma,\mathbf{w}})^j (1 - P_{\sigma,\mathbf{w}})^{d+1-j} \mathbf{w}(j) \right],$$

completing the proof. □

The proof of Theorem 1 does not rely on any particular referral cost structure. The constraints (a), (b), and (c) are necessary under any cost structure as long as high-degree players have the ability to mimic low-degree players. The construction of $\mathbf{w}$ to implement a sequence $\{(r_d, I_d)\}$ works without modification for any monotonic cost function.

Even with this relatively straightforward characterization of implementable sequences, the firm's problem as given by (9) is potentially difficult because the constraints are coupled through the consumer equilibrium. Rather than solving this problem directly, we consider first a relaxed problem in which we only impose the feasibility and individual rationality constraints. This allows us to decouple each element of the sum, finding separately for each degree d the pair (r_d, I_d) that solves

$$\begin{aligned} \max_{r_d,\, I_d} \quad & (1 - F(p - I_d))(p - I_d + r_d(y - \delta)) \\ \text{s.t.} \quad & I_d \ge 0 \\ & r_d \in \{0, 1, \ldots, d\}. \end{aligned} \tag{10}$$

We can usefully view the firm's choice as a two-dimensional mechanism design problem. Similar to a classical adverse selection problem, the firm would like to discriminate between consumers across both valuations and degrees. The relaxation above is equivalent to ignoring incentive constraints with regards to the focal consumer's degree, and simply solving a collection of standard pricing problems, one for each possible degree. The following theorem shows that the solution of the relaxed problem (10) satisfies our incentive compatibility constraints, and thus produces a solution for the original problem. In essence, we show that the two-dimensional problem decomposes into a collection of one-dimensional problems. We can do this because of the natural ordering of the action sets as a function of degree: the firm wishes to offer larger discounts to high-degree consumers, and low-degree consumers cannot mimic high-degree ones.

Theorem 2. *There is a solution to problem* (10) *in which* $r_d = d$ *for all* d, *and the* I_d *are nondecreasing in* d. *Consequently, this solution also optimizes problem* (9). *If we further have that*

$$(p - I) - \frac{1 - F(p - I)}{f(p - I)} \quad \text{is decreasing in } I, \tag{11}$$

then the optimal sequence $\{(r_d, I_d)\}_{d\in\mathbb{N}}$ *is unique.*[6]

Proof. Since $y - \delta > 0$, the firm should clearly pay for all available referrals, so we take $r_k = k$ for each k. The optimal I_k solves

$$\max_{I_d} (1 - F(p - I_d))(p - I_d + d(y - \delta)),$$

subject to $0 \le I_d \le p$. The first derivative of the objective is

$$f(p - I_d)(p - I_d + d(y - \delta)) - (1 - F(p - I_d)).$$

Observe that for fixed I_d this derivative is increasing in d, so the value of I_d that attains the maximum is nondecreasing in d. If this derivative is positive

[6] Note that Equation (11) is the classical regularity of the virtual valuations from Myerson's seminal paper on auction design (Myerson 1981).

(negative) for all $I_d \in [0, p]$, the unique solution is $I_d = p$ ($I_d = 0$). If Equation (11) holds, and we have an interior optimum, the solution is the unique value

$$I_d = \sup\left\{x\colon \frac{1-F(p-x)}{f(p-x)} - (p-x) \le d(y-\delta)\right\}. \quad \square \tag{12}$$

Theorem 2 demonstrates that implementability does not meaningfully restrict the referral and discount pairs we can offer to consumers. We can achieve perfect price discrimination along consumer degrees and decouple the optimization for each part of the sum. As in Theorem 1, this result is robust to many alternative cost specifications. As long as a consumer's incremental cost of making one more referral is always less than y, the argument goes through without modification: the firm wishes to pay for all possible referrals and provides discounts that are increasing in degree. If consumers face a fixed cost for making any positive number of referrals on top of the linear cost δ per referral, it may no longer be optimal to pay for referrals from low-degree consumers. We can describe the solution to the corresponding relaxed optimization problem using a threshold degree d: the firm pays for no referrals from consumers with degree less than d and pays for all referrals from consumers with degree at least d, offering discounts that increase in degree. Such a policy is implementable using the same argument in Theorem 1.

We can better understand the optimal policy if we consider the effective price $p - I_d$ consumers of degree d face. Take $d = 0$ in (12), and note that

$$p - I_0 - \frac{1-F(p-I_0)}{f(p-I_0)} = 0$$

corresponds to the standard monopoly price solution (i.e., virtual valuation is equal to zero). For every additional neighbor a consumer has, the value of I_d increases, indicating a lower effective price.[7] The more neighbors a consumer has, the greater the incentive to trade off immediate profits for additional referrals. Moreover, Equation (12) tells us how to divide the added surplus y between the consumer and the firm: the consumer pockets the direct change in I_d, and the firm earns the corresponding change in the ratio $(1-F(p-I_d))/f(p-I_d)$.

The practicality of optimal referral payments depends on the ease with which we can specify them and the ease with which consumers can respond to them. In general, the optimal payment function faces challenges on both counts because it will be nonmonotonic in the number of successful referrals. Perhaps surprisingly, we might want to offer a consumer \$20 for three referrals, but only \$15 for four.

[7] Note that these optimal effective prices are subject to boundary constraints. If p is below the monopoly price, then low-degree consumers may receive no discount at all; sufficiently high-degree consumers will always receive a full discount, facing an effective price of zero.

Figure 1 **(Color online) Optimal Expected Discount I as a Function of Degree, and the Corresponding Nonmonotone Payment Function w**

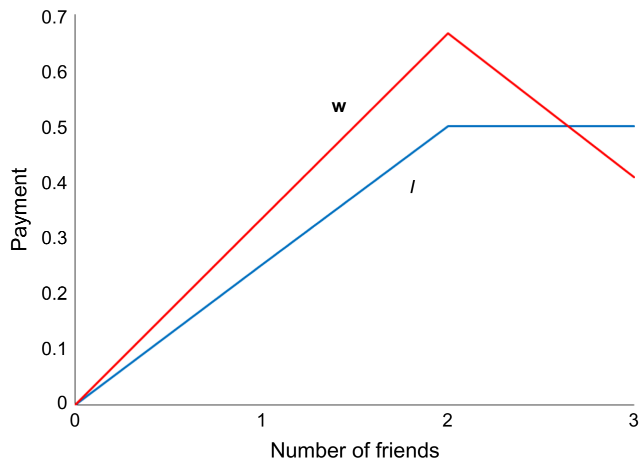

To understand why this happens, consider a simple example in which F is uniform, all consumers have at most three neighbors, the cost of making referrals δ is zero, and the firm's price and value for advertising are $p = y = \frac{1}{2}$. Figure 1 shows the optimal discounts of $I_0 = 0$, $I_1 = \frac{1}{4}$, and $I_2 = I_3 = \frac{1}{2}$. Suppose the degree distribution is such that the corresponding neighbor purchase probability is $P_{\sigma,\mathbf{w}} = \frac{3}{4}$. We compute the corresponding payment function $\mathbf{w}$ as $\mathbf{w}(1) = \frac{1}{3}$, $\mathbf{w}(2) = \frac{2}{3}$, and $\mathbf{w}(3) = 11/27 < \frac{2}{3}$. Because of the randomness in neighbor purchase decisions, the discount a consumer receives is a weighted average over values of the payment function $\mathbf{w}$. A consumer with three neighbors has a higher chance of two successes than a consumer with only two neighbors. To equate I_2 and I_3, we must reduce the benefit from the third successful referral to balance out this effect.

Nonmonotonic payments are a generic feature of the optimal policy whenever the cost δ of making referrals is small. Since neighbors respond stochastically to referrals, the optimal expected payment to a consumer of degree d is a smoothed version of the payment function $\mathbf{w}$. For sufficiently high d, the optimal discount I_d flattens out, and this creates a kink in the optimal expected payment $I_d + d\delta$ as a function of d. For a smoothed version of $\mathbf{w}$ to exhibit this kink, the payment function $\mathbf{w}$ must exhibit a much sharper change. If δ is sufficiently small, this entails a decrease in the payment.

Proposition 3. *Suppose the optimal sequence of referral and discount pairs is such that $0 < P_\sigma < 1$. Holding the net referral value $y - \delta$ fixed, there exists $\bar{\delta} \ge \underline{\delta} > 0$ such that whenever $\delta < \underline{\delta}$, the optimal payment function $\mathbf{w}$ is* *nonmonotonic, and whenever $\delta > \bar{\delta}$, the optimal payment function* $\mathbf{w}$ *is monotonic.*

Proof. See the online appendix (available as supplemental material at http://dx.doi.org/10.1287/mnsc.2016.2476). □

If δ is very small, there is a point at which the optimal expected payment to consumers flattens out. This always occurs when the optimal effective price reaches the point at which all consumers would buy.[8] To achieve this with a smoothed version of $\mathbf{w}$ requires a decrease in $\mathbf{w}$.

In practice, firms do not employ nonmonotonic payments. Not only are these payment functions difficult to compute, but real consumers are unlikely to respond optimally. To address this shortcoming, we study approximations based on empirically common referral payment structures. We consider how well linear and threshold payment functions can approximate the optimal scheme, and we look at how the structure of the network and the cost of making referrals affects which type of referral program performs better.

5. Approximate Incentives

A linear payment function $\mathbf{w}_L$ is one that gives the consumer a fixed payment for each successful referral, i.e., $\mathbf{w}_L(N) = aN$ for some $a \geq 0$. Using a linear $\mathbf{w}_L$, we can implement any sequence of pairs $\{(r_d, I_d)\}$ with $r_d = d$ and $I_d = dI$ for some $I \geq 0$. A threshold payment function $\mathbf{w}_T$ makes a fixed payment b once some threshold number of successful referrals is met. We have $\mathbf{w}_T(N) = b$ if and only if $N \geq \tau$, and $\mathbf{w}_T(N) = 0$ otherwise, for some $b \geq 0$ and $\tau \geq 0$. The set of implementable pairs is less straightforward to describe for threshold payments because it depends on changes in the probability of crossing a given threshold. In either case, the restriction on the set of pairs $\{(r_d, I_d)\}$ means we lose the ability to decouple the summands in our optimization problem (9): the degree distribution now matters.

Our analysis focuses on worst-case degree distributions. Clearly, some degree distributions entail no approximation loss—for instance, if all consumers have the same number of neighbors. A worst-case analysis helps us understand how robust different types of referral programs are. The optimal profit, and hence the potential approximation loss, scales with the average degree in the network $\mu \equiv \mathbb{E}_{D\sim G}[D]$, holding y fixed. To make meaningful comparisons, we consider worst-case performance over distributions with a given average degree. We look at how the worst-case approximation loss scales with the average degree in the network.

[8] We assume here that all consumers would buy at a price of zero, but the same nonmonotonicity result would obtain as long as there is some (possibly negative) price at which all consumers buy.

Given a degree distribution G, let $\pi^*(G)$ denote the firm's profit using the optimal referral payments, let $\pi_L(G)$ denote the firm's profit using the best linear payment function, and let $\pi_T(G)$ denote the firm's profit using the best threshold payment function. We define the *worst-case loss* from linear payments in a network with average degree μ as

$$L_L(\mu) = \sup_{G:\, \mathbb{E}_{D\sim G}[D]=\mu} \pi^*(G) - \pi_L(G),$$

and we define the worst-case loss from threshold payments in a network with average degree μ as

$$L_T(\mu) = \sup_{G:\, \mathbb{E}_{D\sim G}[D]=\mu} \pi^*(G) - \pi_T(G).$$

The optimal profit scales linearly in μ as the value of referrals becomes the main benefit the firm obtains. Approximation losses from using linear payments grow much more slowly, scaling with $\sqrt{\mu}$, while those from using threshold incentives grow linearly. This indicates that linear referral payments are typically preferable when μ is large.

Theorem 3. *Approximation losses from using linear payments $L_L(\mu)$ scale according to $\sqrt{\mu}$. That is, there are constants $\underline{a}$, $\bar{a}$, and a' such that*

$$\underline{a}\sqrt{\mu} \leq L_L(\mu) \leq \bar{a}\sqrt{\mu} \quad \text{for all } \mu \geq a'.$$

If $\delta > 0$, approximation losses from using threshold payments $L_T(\mu)$ scale linearly with μ. That is, there are constants $\underline{b}$, $\bar{b}$, and b' such that

$$\underline{b}\mu \leq L_T(\mu) \leq \bar{b}\mu \quad \text{for all } \mu \geq b'.$$

Proof. Consider the linear approximation first. To establish the upper bound, note that for sufficiently large average degree μ the optimal policy for all consumers with degree d higher than $\sqrt{\mu}$ is to offer $I_d = p$, yielding profit of dy per consumer. For consumers of degree less than $\sqrt{\mu}$, the optimal profit is no more than $\sqrt{\mu}(y-\delta)$ per consumer since profits per consumer are increasing in degree.

Suppose we choose a linear discount policy $I'_d = dI'$ with $I' = p/\sqrt{\mu}$ for all d. For the consumers with degree above $\sqrt{\mu}$, the loss due to overpayment is at most $p/\sqrt{\mu}$ per referral. We can bound the loss from overpayments by

$$\mathbb{E}[I'_d] = \frac{p\,\mathbb{E}_{D\sim G}[D]}{\sqrt{\mu}} = p\sqrt{\mu}.$$

Likewise, since the total profit from consumers with degree lower than $\sqrt{\mu}$ is bounded by $\sqrt{\mu}(y-\delta)$, this is a bound on the loss as well because profits from a consumer are never negative. This proves the upper bound.

For the lower bound, consider the degree distribution G_μ that assigns degree $\lfloor\sqrt{\mu}\rfloor$ with probability $\mu/(2\mu-\lfloor\sqrt{\mu}\rfloor$ and degree 2μ with probability $(\mu-\lfloor\sqrt{\mu}\rfloor)/(2\mu-\lfloor\sqrt{\mu}\rfloor)$. This distribution has expectation μ, and for sufficiently large μ, the optimal policy sets $I_{\lfloor\sqrt{\mu}\rfloor}=I_{2\mu}=p$.

Let $I'_d = dI'$ denote a linear policy, and consider two cases. Suppose first that $I' \le p/(2\lfloor\sqrt{\mu}\rfloor)$. Then low-degree consumers receive a discount of no more than $p/2$, and lost sales coupled with lost referrals lead to approximation losses of at least $F(p/2)\lfloor\sqrt{\mu}\rfloor(y-\delta)$ from each low-degree consumer. At least half of consumers have low degrees, so half of this is a lower bound on the approximation loss. Alternatively, suppose $I' > p/(2\lfloor\sqrt{\mu}\rfloor)$. For large μ, essentially all referrals come from high-degree consumers, and essentially all of this discount is overpayment. For large μ, the overpayment losses converge to $(\mu-\lfloor\sqrt{\mu}\rfloor)p/(2\lfloor\sqrt{\mu}\rfloor)$. In both cases, the worst-case loss is at least a constant times $\sqrt{\mu}$, proving the claim.

Now consider the threshold approximation. The upper bound is trivial since optimal profits scale linearly with μ. For the lower bound, consider for each μ the degree distribution that places probability $\frac{1}{2}$ on $\mu/2$ and probability $\frac{1}{2}$ on $3\mu/2$.

Suppose that for some large μ, the threshold τ is larger than $\mu/2$. None of the low-degree consumers make any referrals, which leads to losses of $(\mu/4)(y-\delta)$. Alternatively, suppose that $\tau \le \mu/2$. To avoid similar losses, we must incentivize the high-degree consumers to refer at least μ of their neighbors. From Hoeffding's inequality, for large μ, the probability of *not* hitting the threshold number of successes τ after μ referrals is on the order of $e^{-\mu/2}$. If adding one more referral is to increase the probability of hitting τ enough to compensate for the cost δ, the threshold payment must be very large, growing exponentially in μ. This implies overpayments to high-degree consumers that grow exponentially in μ. Hence, for large μ the optimal threshold payment scheme will choose a high threshold τ, yielding linear losses because the low-degree consumers do not refer any neighbors. □

The proof of Theorem 3 sheds light on why linear and threshold payments perform well or poorly. Linear payments suffer when the degree distribution is heavy tailed. The discount that a consumer receives is proportional to how many neighbors she has. When most consumers have low degrees, but most referrals come from those with high degrees, linear payments face a difficult trade-off. Providing the appropriate discount to low-degree consumers entails large overpayments to high-degree ones, while giving an appropriate discount to high-degree consumers entails fewer purchases from low-degree ones. The overpayment problem arises regardless of the model parameters, as long as there is a limit to how much the firm should optimally pay a consumer to take the product.

This asymptotic result is robust to many modifications we might make to the model or the payment function. In the previous section, we discussed alternative specifications for consumer referral costs. What matters for the asymptotic performance of these policies, and the proof provided, is the marginal referral cost when making a large number of referrals. For instance, we can apply exactly the same argument in the case with a fixed cost for making any positive number of referrals in addition to the linear cost δ per referral. Perhaps surprisingly, adding a payout cap to a linear referral policy yields no improvement in our asymptotic performance bounds. Although the firm can use a cap to prevent overpayment to very high-degree consumers, these consumers will respond by not referring all of their neighbors. Using an argument similar to that for the threshold payment function result, we find the losses from foregone referrals are comparable to those from overpayment.

If we can constrain the tail of the degree distribution, the losses from using linear payments scale more slowly. Given a constant $C>0$ and an exponent $\alpha\in[0,1)$, consider the collection of all distributions $\mathcal{S}_C^\alpha(\mu)$ with average degree μ such that the tail above the mean satisfies

$$\sum_{d\ge\mu}(d-\mu)g_d \le C\mu^\alpha.$$

The exponent α is the crucial part of this constraint; with larger α, the permitted length of the tail scales more rapidly with the average degree. We can interpret this as a restriction on how large the mean is conditional on an above-average degree. The case $\alpha=0$ implies the conditional mean is no larger than μ plus a constant, while an $\alpha>0$ allows the difference between μ and the conditional mean to grow. We define the (C,α) *worst-case loss* from using linear payments as

$$L_C^\alpha(\mu) = \sup_{G\in\mathcal{S}_C^\alpha(\mu)} \pi^*(G) - \pi_L(G).$$

Proposition 4. *Fixing C and α, the (C,α) worst-case loss from using linear payments scales no faster than $\mu^{\alpha/2}$. That is, there exist constants $\bar{a}$ and a' such that*

$$L_C^\alpha(\mu) \le \bar{a}\mu^{\alpha/2} \quad \textit{for all } \mu \ge a'.$$

Proof. See the online appendix. □

Though threshold payments have worse asymptotic properties than linear payments, the conditions that hamper their performance are distinct. Having just one payment limits the flexibility of these schemes to compensate the marginal cost of making referrals. If these costs are significant, then a threshold payment can only target consumers within a narrow range of degrees. Those with lower degrees cannot benefit from

Table 1 Percentage of Optimal Profit

	Poisson (%)	Geometric (%)	Power law (%)
Linear payment	100	99.1	90.8
Threshold payment	99.4	97.9	98.5

the threshold payment since they will never hit the threshold. Those with higher degrees still receive an appropriate discount, but they refrain from referring all of their neighbors because the marginal impact on the probability of hitting the threshold is too small to compensate for the additional cost.

In networks with relatively low average degrees, these different strengths and weaknesses can make either linear or threshold payments preferable. To see this, consider an example comparing the performance of linear and threshold payments across different networks. For these calculations, we assume consumer valuations are uniform, the price of the product is $p = 0.5$, the cost of making referrals is zero, and the value of referrals to the firm is $y = 0.08$. We compute optimal linear and threshold payments for three different degree distributions: a Poisson distribution, a geometric distribution, and a power-law distribution with exponent 2.1. In each case, the average degree in the network is approximately 3.6. Table 1 reports the percentage of the optimal profit we obtain in each case.

Performance is quite good across the board; having a referral program, even a very simple one, is far better than having none. We can see a clear decline in the performance of linear payments as the tail of the degree distribution becomes fatter. Moving from a Poisson to a geometric distribution costs 1% of profits, and moving to the much fatter-tailed power-law distribution costs an additional 8%. Threshold payments are more consistent across the distributions. For the Poisson and the geometric distribution, threshold payments perform slightly but noticeably worse than linear ones. For the power-law distribution, threshold payments offer a significant improvement. Note the absence of referral costs is important to this comparison. If we increase δ while holding the net referral value y fixed, the performance of linear payments will remain unchanged, but threshold payments will suffer.

Linear and threshold payments have complementary strengths. Linear payments can effectively compensate the cost of making referrals, while threshold payments can provide a discount to high-degree consumers that remains bounded. A combination of these two can lead to much better performance. Consider the class of linear payment functions with a "bonus" for crossing some threshold. These are payment functions $\mathbf{w}_B$ we can express as

$$\mathbf{w}_B(N) = \mathbf{w}_L(N) + \mathbf{w}_T(N),$$

where $\mathbf{w}_L$ is linear and $\mathbf{w}_T$ is a threshold payment function. Let $L_B(\mu)$ denote the worst-case loss using such payment functions when the average degree in the network is μ. With this slight increase in complexity, we can achieve a constant bound on losses.

Theorem 4. *The worst-case loss using linear payments with a bonus is bounded by a constant. There exists* $C > 0$ *such that*

$$L_B(\mu) \le C \quad \textit{for all } \mu.$$

Proof. We prove a stronger result, showing that a large class of policies achieves constant loss. Choose the linear component $\mathbf{w}_L$ to exactly compensate the social cost δ per referral. This reduces our problem to showing that a threshold payment achieves constant loss when the social cost of referrals is zero.

In the specification of the threshold payment function $\mathbf{w}_T$, choose any $c > p$ and any finite K. Since $P_{\sigma,\mathbf{w}} \ge 1 - F(p) > 0$, there is some threshold degree $\bar{d}$ such that for any purchasing consumer with $d \ge \bar{d}$, the probability that at least K neighbors purchase in response to a referral is at least p/c. These consumers receive a full discount and overpayment bounded by $c - p$. Losses from consumers with a degree lower than $\bar{d}$, or for whom the optimal policy does not offer a full discount, are bounded by a constant because there are finitely many such degrees. Losses from higher degree consumers are bounded by $c - p$. □

To avoid significant losses from the highest degree consumers, we need to offer a full discount and we need to compensate for each referral. Choosing a linear component that exactly compensates for the social cost δ ensures that any purchasing consumer makes all possible referrals, and we can adjust the threshold payment to provide a sufficiently large discount for the high-degree consumers. The size of the threshold payment gives us a fixed bound on any over payments, yielding the constant loss bound in Theorem 4.

An implication of Theorem 4 is that, in the absence of any cost for making referrals, threshold payments alone can achieve constant loss.

Corollary 1. *If* $\delta = 0$, *approximation losses from using threshold payments* $L_T(\mu)$ *are bounded by a constant.*

Together with our earlier findings, this suggests that threshold payments can outperform linear ones in some settings. Whenever we have the following:

(a) the degree distribution has a fat tail, so linear payments perform relatively poorly, and

(b) the marginal cost of referrals is low, so threshold payments perform relatively well,

threshold payments may be the best option among the simplest payment schemes. Notably, conditions (a) and (b) are characteristic of many social-media settings: the degree distribution in many social networks follows a power law (e.g., Ugander et al. 2012), and referrals can often be made en masse using a single post.

6. Extensions

6.1. Differentiating Conversions from Nonconversions

Although there is certainly value to a company from referrals beyond any immediate conversions, it makes sense that a firm would place a higher value on referrals that convert than referrals that do not. We can extend our model to account for this, supposing that the firm values unconverted referrals at $\underline{y}$ and converted referrals at $\bar{y}$ with $\bar{y} > \underline{y} > \delta$. The expected referral value is then endogeneously determined in equilibrium as $y = P_{\sigma,\mathbf{w}}\bar{y} + (1 - P_{\sigma,\mathbf{w}})\underline{y}$.

Although this adds an extra dimension to the firm's optimization problem, the essence of our results remains unchanged. The firm's problem in Equation (9) no longer decouples along player degrees since y depends on the strategies of all players, but we can still solve a relaxed version of the problem that only imposes individual rationality constraints. The solution of the relaxed problem will still be incentive compatible, so we can find a payment function to implement the optimal sequence of pairs $\{(r_d, I_d)\}$. Knowing that $\delta < \underline{y} \leq y \leq \bar{y}$ is enough to reproduce the asymptotic bounds of Section 5 using the same arguments.

6.2. Optimal Pricing

If the firm jointly optimizes the product price and the referral program, how should the price be set? If there are no Veblen effects and we can use the optimal payment function, then the answer is obvious: choose the standard monopoly price. For consumers with no friends to refer, the monopoly price maximizes profits, and the optimal payment function allows us to give the best possible sequence of discounts from that reference point. If all consumers have at least one neighbor, then the optimal price is nonunique: anything higher than the monopoly price is optimal. This is because we can implement an arbitrary nondecreasing sequence of discounts, so the effective prices we charge consumers will not change with the price p as long as it is above the monopoly price.

If we are using an approximate payment function, the optimal price is less clear. Since the payment function depends on fine details of the degree distribution, the optimal price will as well. At least in the case of linear incentives, we can make a general claim about the optimal price: the standard monopoly price serves as an upper bound.

Proposition 5. *Suppose virtual valuations $v - (1 - F(v))/f(v)$ are increasing and (weakly) convex. Then if we are using a linear payment function* $\mathbf{w}$, *the optimal price is no larger than the standard monopoly price.*

Proof. First, suppose the valuation distribution is uniform. Figure 2 shows the optimal discount as a function of degree as well as a linear approximation. Decreasing the price p shifts the optimal discount curve downward, and the curve hits the point $(0, 0)$ precisely when p is the standard monopoly price. Note that for p equal or greater than the monopoly price, any linear approximation intersects the curve at exactly one point. If we shift the curve downward we can always obtain a strictly better approximation by flattening the line to maintain the same point of intersection: the resulting line is closer to the optimal incentive at every point. This same graphical argument applies whenever the optimal discount curve is weakly concave. Since we have assumed virtual valuations are convex, Equation (12) implies that the optimal discount curve is concave, yielding the result. □

Figure 2 (Color online) If p is Above the Monopoly Price, Reducing p Results in a Better Approximation

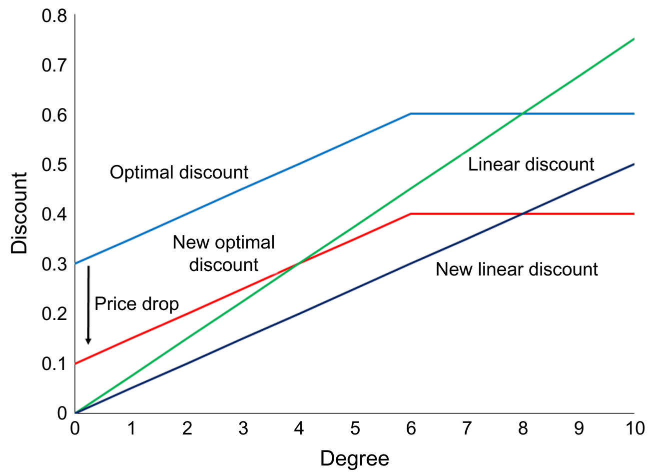

6.3. Two-Way Incentives

Many referral programs make use of two-sided incentives: in addition to giving the referring customer a reward, the referred friend receives a discount. Why should the firm offer such incentives? In our framework, there is little reason to include this feature since we already have as much flexibility as we could want to adjust the expected referral reward. The only effect would be to reduce revenues from referred individuals, which would imply the firm should not offer this type of incentive. However, given the prevalence of such programs, this answer is unsatisfying.

One possible explanation is that the additional incentive may reduce social costs. Part of the cost of making referrals is psychological, stemming from a desire not to annoy one's friends or appear self-serving. If making a referral confers a benefit to the friend receiving it, these psychological costs may shrink. We could extend our model to make the social cost δ a function of the discount provided to the friend receiving the referral. The firm would then jointly optimize the referral reward function and the net referral value by manipulating the social cost. Once the social cost is fixed by choosing the discount for referred friends, the problem reduces to the one we study, so our principal findings remain unchanged.

6.4. Additional Network Information

Individuals often have some information about how well connected their friends are. In this subsection, we explore the implications of consumers having signals about which of their friends are well connected. Suppose a consumer receives a binary signal about each neighbor, indicating whether the neighbor has degree greater than some threshold τ. The consumer knows that her neighbors receive similar information about their neighbors, and that neighbors of neighbors have degrees drawn independently according to distribution G'' as before. The firm will naturally place a higher value on referrals sent to higher degree neighbors—we write y^+ for the expected value of a referral to a consumer with $d > \tau$, and y^- for the value of a referral to a consumer with $d \le \tau$. We assume $y^+ > y^- \ge \delta$. We consider first how the consumer equilibrium changes before looking at the effects of this information on the firm's design problem.

We now characterize a consumer using her valuation v, her degree d, and the number of high signals h she receives. Fixing the payment function $\mathbf{w}$, a symmetric strategy profile σ maps triples (v, d, h) into actions that specify whether the consumer buys $b \in \{0, 1\}$, the number of low-degree referrals to send $r^- \in \{0, 1, \ldots, d-h\}$, and the number of high-degree referrals to send $r^+ \in \{0, 1, \ldots, h\}$. A consumer's best reply then depends on a pair of neighbor purchase probabilities $\mathbf{P} = (P^-, P^+)$ for the corresponding types of neighbors. We can define a pair of best reply neighbor purchase probabilities $(\phi^-(\mathbf{P}), \phi^+(\mathbf{P}))$ analogously to Equation (7). Given a strategy profile σ, write P_σ^- and P_σ^+ for the corresponding probabilities that a low-degree and a high-degree neighbor following σ will purchase in response to a referral.

Proposition 6. *A symmetric pure strategy Bayesian equilibrium of the consumer game exists. In any such equilibrium σ, we must have $P_\sigma^+ \ge P_\sigma^-$.*

Proof. See the online appendix. □

Although the consumer problem is very similar to the one we considered in Section 3, with additional network information, the firm can no longer decouple its optimization for each possible consumer degree. To understand why, consider the following simple example. Let v be uniform on $[0, 1]$, let $p = \frac{1}{2}$, and suppose G takes the value 1 with probability $\frac{2}{3}$ and 2 with probability $\frac{1}{3}$, so that G'' is equally likely to take the values 0 and 1. We suppose that the signal allows a consumer to identify whether a neighbor has degree 0 or 1.

Consider first the problem of designing a policy that maximizes profits obtained from a focal consumer with degree 1. This focal consumer can have either a low-degree or a high-degree neighbor ($d = 0$ or 1). We can describe this consumer's behavior using a referral and discount pair for each case: write (I_1^-, r_1^-) for the pair corresponding to a consumer with a low-degree neighbor and (I_1^+, r_1^+) for that corresponding to a consumer with a high-degree neighbor. Any pairs that are implemented in equilibrium must satisfy the following individual rationality and incentive compatibility constraints:

- $I_1^- > 0$ if and only if $r_1^- = 1$; $I_1^+ > 0$ if and only if $r_1^+ = 1$;
- $r_1^+ \ge r_1^-$ and $I_1^+ \ge I_1^-$;
- $I_1^- + \delta \ge (1/(1 + I_1^- + I_1^+))(I_1^+ + \delta)$.

The first two constraints are immediate from the problem definition and that high-degree neighbors are more likely to purchase than low-degree ones (Proposition 6). The last constraint arises because a consumer with a low-degree neighbor could mimic one with a high-degree neighbor. The payment $\mathbf{w}(1)$ for a successful referral is the same in both cases, so the difference in the expected payment to each type is limited by the probability of a successful referral: a low-degree neighbor will buy with probability $\frac{1}{2}$, and a high-degree neighbor will buy with probability $\frac{1}{2} + (I_1^- + I_1^+)/2$.

Because of this last constraint, we cannot decouple the optimization for each type of consumer. Depending on the problem parameters, we might forego referrals to low-degree neighbors to avoid giving too much surplus to consumers with high-degree neighbors, or we might allow those with high-degree neighbors to get extra surplus in order to obtain all possible referrals. For a consumer with more neighbors, we face additional nontrivial incentive compatibility constraints. Therefore, the technique we proposed in Section 4, which involved solving a two-dimensional mechanism design problem by solving a family of one-dimensional relaxations, does not apply to the three-dimensional mechanism design problem that emerges when customers have additional network information.

Given this challenge, we can ask, how robust is our solution to the original problem? Put differently, if the firm ignores that consumers have information about one another, how does the corresponding referral policy perform? The second claim in Proposition 6 shows that consumer incentives for making referrals are largely aligned with the firm's interests: high-degree neighbors are more valuable to the firm, and these are precisely the neighbors most likely to respond to a referral and generate payments for the consumer. Consequently, we might expect this information to benefit the firm.

We can prove for a class of problem instances that this is indeed the case. For simplicity, we make the following assumptions.

Assumption 1. *The valuation distribution v is uniform on the interval* $[0, 1]$, *the price is the standard monopoly price* $p = \frac{1}{2}$, *and no consumer has degree higher than* $1/(y^+ - \delta)$.[9] *Let* $y = (y^- + y^+)/2$ *denote the expected value of a referral to a random neighbor when no information is observed, and assume that* $y \geq 3\delta$. *The threshold* τ *is such that a random neighbor generates low and high signals with equal probability.*

If consumers receive no information about their neighbors, the firm's optimal policy is a linear payment function that provides an effective discount of $(y-\delta)/2$ per neighbor. Given the degree distribution, these effective discounts correspond to a particular neighbor purchase probability P^*, and the optimal payment function is

$$\mathbf{w}(N) = \frac{(y+\delta)N}{2P^*}. \tag{13}$$

We now consider what happens if the firm applies this policy when consumers do observe signals about their neighbors.

Proposition 7. *Suppose Assumption* 1 *holds. Suppose also the firm employs the optimal policy under the assumption of no consumer information about their neighbors, as described in Equation* (13). *Then, the firm will earn higher profit if consumers receive signals about their neighbors' degrees than if they do not.*

Proof. See the online appendix. □

Proposition 7 suggests that the firm will often benefit more from referrals when consumers have additional information about one another. Moreover, the firm can employ the optimal policy for the simpler model without this information, and the estimated optimal profit serves as a lower bound on the profit obtained. Interestingly, the increase in profit is not the result of a focal consumer becoming more likely to purchase. The key driver of this result is that consumers are carrying out better referral targeting on the firm's behalf: consumers preferentially refer high-degree neighbors.

7. Final Remarks

Start-up companies increasingly rely on referrals to launch their businesses, and many of the great success stories in recent memory have been driven by customer referral programs. The rise of social media makes it easier for companies to track referrals and provide associated payments, greatly reducing the costs of implementing these programs relative to traditional advertising. Social media also makes it easier for consumers to refer their friends; emailing a link or making a Facebook post requires far less effort than bringing up a product in conversation. These features will ensure a continuing role for referral programs in firms' advertising mix.

Our decomposition of referral payments highlights two distinct considerations when we design a referral program: we need to compensate consumers for making referrals, and we need to discount the product for those who can bring in many referrals. Since there is a limit to the discount we wish to offer, the latter piece introduces a kink in the optimal discount as a function of how many neighbors a consumer has. With no restrictions on our referral payment function, this kink can lead to unwieldy optimal payments that are nonmonotonic in the number of successful referrals.

The same decomposition can guide our thinking when we design simple payment schemes to approximate the optimal one. In our model, linear payments can effectively compensate consumers for making referrals, and threshold payments can provide a discount to high-degree consumers while limiting over payments. As a result, each type of referral program has different strengths and weaknesses depending on the structure of the social network and the significance of referral costs. Linear payments are more flexible in general and will typically outperform threshold payments. However, the combination of a fat-tailed degree distribution with low marginal referral costs, which is characteristic of many social-media settings, favors threshold payments.

If we can implement a slightly more complex payment scheme, linear payments with a bonus may offer a substantially better option. These functions combine features of both linear and threshold payments. We have two degrees of freedom, one that we can tailor to compensate consumers for making referrals and one that we can use to appropriately discount the product for high-degree consumers. Regardless of the degree distribution or the cost of referrals, this type of referral program can offer a good approximation to the optimal profit.

There are important limitations to our analysis that present opportunities for future research. Although we have included several extensions looking at the robustness of our results, we have certainly not exhausted all possibilities. One important phenomenon not yet discussed is homophily, the tendency for individuals to link to similar others. This suggests that friends should not have independent valuations for a product; they likely hold similar valuations. Although addressing the effects of homophily on referral program design is beyond the scope of our work, our model provides a useful baseline from which to explore this and related extensions.

The firm in our model has limited information about the network of consumers, knowing only the degree distribution. With access to social-media data, it is often

[9] This last condition ensures it is never optimal to fully discount the product, thus avoiding issues that arise because of nonmonotone optimal payments.

possible to learn more about the network position of a consumer, and it is possible to target offers to specific individuals. With this additional information, a firm will wish to discriminate between consumers using more sophisticated measures of centrality. Studying how optimal referral payments change when the firm has access to this information offers another set of extensions to explore.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/mnsc.2016.2476.

Acknowledgments

The authors are grateful to the department editor, the associate editor, and the three reviewers for their comments that helped improve the initial version of this paper. The authors thank IBM (Open Collaborative Research Award) for its support of this research project.

References

Aral S, Walker D (2011) Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Sci.* 57(9):1623–1639.

Aral S, Muchnik L, Sundararajan A (2013) Engineering social contagions: Optimal network seeding in the presence of homophily. *Network Sci.* 1(2):125–153.

Ballester C, Calvó-Armengol A, Zenou Y (2006) Who's who in networks. Wanted: The key player. *Econometrica* 74(5):1403–1417.

Biyalogorsky E, Gerstner E, Libai B (2001) Customer referral management: Optimal reward programs. *Marketing Sci.* 20(1):82–95.

Bloch F, Quérou N (2013) Pricing in social networks. *Games Econom. Behav.* 80:243–261.

Burke M, Kraut R (2014) Growing closer on Facebook: Changes in tie strength through social network site use. *Proc. SIGCHI Conf. Human Factors Comput. Systems (CHI2014)* (ACM, New York), 4187–4196.

Buttle FA (1998) Word of mouth: Understanding and managing referral marketing. *J. Strategic Marketing* 6(3):241–254.

Campbell A (2013) Word-of-mouth communication and percolation in social networks. *Amer. Econom. Rev.* 103(6):2466–2498.

Campbell J (2015) Localized price promotions as a quality signal in a publicly observable network. *Quant. Market Econom.* 13(1): 27–57.

Candogan O, Bimpikis K, Ozdaglar A (2012) Optimal pricing in networks with externalities. *Oper. Res.* 60(4):883–905.

Chatterjee K, Dutta B (2014) Credibility and strategic learning in networks. Working paper, Pennsylvania State University, State College.

Chu S-C, Kim Y (2011) Determinants of consumer engagement in electronic word-of-mouth (eWOM) in social network sites. *Internat. J. Advertising* 30(1):47–75.

Cohen MC, Harsha P (2015) Designing price incentives in a network with social interactions. Working paper, Google Research, New York.

Galeotti A, Goyal S, Jackson M, Vega-Redondo F, Yariv L (2010) Network games. *Rev. Economic Stud.* 77(1):218–244.

Godes D, Mayzlin D (2004) Using online conversations to study word-of-mouth communcation. *Marketing Sci.* 23(4):545–560.

Goldenberg J, Libai B, Muller E (2001) Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Lett.* 12(3):211–223.

Houston D (2010) Dropbox—Startup lessons learned. Presentation slides, April 23, http://www.slideshare.net/gueste94e4c/dropbox-startup-lessons-learned-3836587.

Jackson M, Yariv L (2007) Diffusion of behavior and equilibrium properties in network games. *Amer. Econom. Rev.* 97(2): 92–98.

Kempe D, Kleinberg J, Tardos É (2003) Maximizing the spread of influence through a social network. *Proc. Ninth ACM SIGKDD Internat. Conf. Knowledge Discovery and Data Mining* (ACM, New York), 137–146.

Kumar V, Petersen A, Leone R (2007) How valuable is word of mouth? *Harvard Bus. Rev.* 85(10):139–144.

Leskovec J, Adamic LA, Huberman BA (2007) The dynamics of viral marketing. *ACM Trans. Web* 1(1):Article 5.

Libai B, Biyalogorsky E, Gerstner E (2003) Setting referral fees in affiliate marketing. *J. Service Res.* 5(4):303–315.

Mobius M, Niehaus P, Rosenblat T (2011) Social networks and consumer demand. Working paper, Microsoft Research, Cambridge, MA.

Myerson R (1981) Optimal auction design. *Math. Oper. Res.* 6(1): 58–73.

Sundararajan A (2007) Local network effects and complex network structure. *B. E. J. Theor. Econom.* 7(1):Article 46.

Ugander J, Karrer B, Backstrom L, Marlow C (2012) The anatomy of the Facebook social graph. Working paper, Stanford University, Stanford, CA.