Complex Contagion in Social Networks: Causal Evidence from a Country-Scale Field Experiment

Jaemin Lee
Department of Sociology
The Chinese University of Hong Kong
jaeminlee@cuhk.edu.hk

David Lazer
Network Science Institute
Northeastern University
d.lazer@northeastern.edu

Christoph Riedl*
D'Amore-McKim School of Business
Northeastern University
c.riedl@northeastern.edu

forthcoming in *Sociological Science*

Abstract

Complex contagion rests on the idea that individuals are more likely to adopt a behavior if they experience social reinforcement from multiple sources. We develop a test for complex contagion, conceptualized as social reinforcement, and then use it to examine whether empirical data from a country-scale randomized controlled viral marketing field experiment show evidence of complex contagion. The experiment uses a peer encouragement design in which individuals were randomly exposed to either one or two friends who were encouraged to share a coupon for a mobile data product. Using three different analytical methods to address the empirical challenges of causal identification, we provide strong support for complex contagion: the contagion process cannot be understood as independent cascades, but rather as a process in which signals from multiple sources amplify each other through synergistic interdependence. We also find social network embeddedness is an important structural moderator that shapes the effectiveness of social reinforcement.

Keywords: Complex Contagion, Social Networks, Social Influence, Embeddedness, Behavior Change

Reproducibility Package Statement: Replication code and data available at Open Science Framework (OSF): https://doi.org/10.17605/OSF.IO/VCU3S. In the empirical section of our paper, we analyze a proprietary dataset that was shared with us under an NDA. Hence, we are not able to share the data publicly. We provide a de-identified, stratified sample of 70,000 individuals (roughly 3% of the full data) that preserves treatment balance and key covariate variation. We also provide code that reproduces the main tables and figures of the paper using the anonymized dataset.

* Corresponding author

Introduction

How does behavior spread? Almost 20 years ago Centola & Macy (2007) introduced the idea that the adoption of some behaviors may benefit from social reinforcement from multiple sources. This idea emerged as one of the most influential and paradigm-shifting theories in the study of social influence (DiMaggio & Garip, 2012). The idea was transformative because it suggested a fundamentally different perspective from the longstanding belief in the structural advantages of random ties—i.e., the traditional "strength of weak ties" theory (Granovetter, 1973). Instead, the new theory (termed complex contagion) posited that the adoption of costly behaviors—those benefitting from socially reinforcing signals from multiple neighbors due to credibility, legitimacy, and complementarity—can benefit from redundant ties found in clustered networks. The fundamental individual-level prediction of complex contagion is that contact with additional peers who have adopted the behavior can amplify spillovers (Centola and Macy 2007; similar to threshold models (Granovetter, 1978).

Convincingly demonstrating whether an observed diffusion process shows signs of complex contagion is inherently challenging because of a subtle but fundamental puzzle: when adoption is probabilistic, any additional exposure—whether from the same source or a different one—naturally increases the chance of adoption. This makes it difficult to isolate whether an observed increase in adoption likelihood stems from true social reinforcement (the synergy between multiple sources) or simply from repeated independent chances to adopt. Scholars have made substantial progress on this challenge by developing rich theoretical models (Centola, 2018; Centola & Macy, 2007; Eckles et al., 2024; Wan et al., 2025) and accumulating empirical evidence using observational data (Aral et al., 2009; Aral & Nicolaides, 2017; Christakis & Fowler, 2007; Ferrali et al., 2020; Karsai et al., 2014; Katona et al., 2011), network-level analyses (Fink et al., 2016; Ugander et al., 2012), controlled experiments (Centola, 2010), and studies of low-threshold behaviors (Bakshy et al., 2012). Yet, directly quantifying the *additional boost* in adoption due to social reinforcement from multiple distinct peers—over and above what probabilistic exposure alone would predict—remains a difficult problem. Our contribution is to address this gap with a method that formalizes and estimates this multibody synergy by leveraging recent advances in higher-order interaction models and hypergraph theory (St-Onge et al., 2021).

To tackle this challenge, we develop a unifying framework for testing higher-order social influence and provide causal evidence of social reinforcement in a realistic setting with existing social ties. We first develop a conceptualization of social reinforcement drawing on insights on contagion on hypergraphs (St-Onge et al., 2021) that account for group structure in exposure and relaxes the assumption of deterministic behavior (Eckles et al., 2024; Wan et al., 2025). This approach can be used to quantify the strength of both positive and negative social reinforcement by comparing it against a baseline of a counterfactual simple contagion process. We then use this approach with data from a field experiment to test whether it shows evidence of complex contagion. Building on novel insights in experimental design and statistics (Athey et al., 2018) we design a country-scale randomized controlled field experiment that manipulated the core element behind complex contagion. Specifically, we designed a peer encouragement experiment (Bradlow, 1998; Eckles et al., 2016) to seed pairs of customers with coupon codes for a mobile data product and encouraged them to adopt and share the coupon code with their mobile-network neighbors in a viral marketing campaign. This peer encouragement design generated exogenous variation among the targeted customers' neighbors so that some had one neighbor who received a coupon while others had two. We analyze the data from our experiment using three different econometric methods to confront empirical challenges of causally identifying social influence: (1) an intent-to-treat analysis, (2) a two-stage least squares (2SLS) analysis, and (3) a computational exact *p*-value method that accounts for network interference.

We find robust evidence of positive social reinforcement across all three methods, providing strong support for the presence of complex contagion: the contagion process cannot be seen as independent cascades but is instead a process in which signals from multiple sources reinforce each other. We also find significant treatment effect heterogeneity, where network embeddedness—the extent to which two individuals have friends in common—is an important structural moderator that shapes the effectiveness of social reinforcement.

Our study makes three contributions to our understanding of complex contagion theory and the spread of behavior more generally. First, our study extends prior models of complex contagion in an integrative model of social influence based on social reinforcement. The earliest models of complex contagion offered a stylized conceptual foundation (Centola & Macy, 2007) that have since been enriched by a substantial body of work. A growing consensus in the literature suggests that probabilistic formulations offer advantages for modeling complex contagion in real-world contexts (Eckles et al., 2024; Wan et al., 2025) and entirely new modeling approaches building on hypergraphs have opened new avenues in studying higher-order phenomena (St-Onge et al., 2021, 2022). This enriches the empirical interpretation of complex contagion, moving from a threshold-based process where adoption *requires* social reinforcement from multiple sources to a probabilistic process in which adoption is *amplified* by social reinforcement in a synergistic way. This approach relies on local interactions and works without the need to characterize network-level diffusion patterns which can fail to detect positive social reinforcement (Wan et al., 2025). Our method builds on prior theoretical arguments by providing a robust empirical test distinguishing higher-order interaction against a lower-order null model with only pairwise interdependence between people (i.e., capturing the difference between repeat exposure to the same person vs. exposure to different people). The contrast with simple contagion remains powerful. This approach offers a principled way of capturing the strength of social reinforcement, giving researchers a canonical reference point to answer research questions around complex contagion. Our method, furthermore, flexibly handles both positive as well as negative (dampening) social reinforcement and can thus powerfully contribute to study of complex coordination (polarization, echo chambers) and (mis)information.

Second, we provide strong causal evidence of positive social reinforcement from a country-scale randomized field experiment to advance our understanding of the diffusion of products and services (Granovetter, 1978; Rogers, 2003; Strang & Soule, 1998; Valente, 1995). Our study draws on a realistic encouragement design—with existing social ties and voluntary interactions—to directly manipulate the core concept behind complex contagion: whether individuals had contact with one versus two sources. This is particularly useful and consequential for researchers and practitioners looking for empirical grounding of the theoretical predictions for behavior change. Despite recent qualifications showing that network clustering is neither necessary nor always optimal to spread behavior (Wan et al. 2025), social reinforcement fundamentally matters. When social reinforcement is sufficiently strong as in our case, groups are more effective seeds than individual hubs (St-Onge et al., 2022). This insight shapes how we should design interventions, e.g., by nudging clusters, seeding groups instead of single influencers, and using peer reinforcement. Our work thus strengthens the consensus in prior work that have variously shown signs of complex contagion in observational data (Aral & Nicolaides, 2017; Fink et al., 2016; Katona et al., 2011; Mønsted et al., 2017), experimental studies using artificial social networks (Centola, 2010), relied on automated notification systems (Aral & Walker, 2014), and those showing effects for low-threshold behavior (Bakshy et al., 2012; see Table A5 for a summary of how our study differs from closely related prior work).

Third, we contribute to our understanding of the role of embeddedness in social reinforcement. Building on prior work that has pointed to stronger social influence along structurally embedded ties (Aral et al., 2009; Bakshy et al., 2012; Katona et al., 2011; Mønsted et al., 2017), we highlight the role of embeddedness as *relational characteristics* of redundant ties that moderates the effect of social reinforcement.

Background

Social Reinforcement in Adoption of New Behavior

In the canonical view of contagion, spreading can occur based on a single contact with an "infected" (or informed) individual. Joe has the flu, interacts with Anne, who then becomes infected. When projected onto behavioral contagion it is plausible that some behaviors benefit from social reinforcement or affirmation from multiple different sources (Centola & Macy, 2007). The literature discusses several social mechanisms that can explain why exposure to multiple different sources of a behavior may amplify behavior spread: credibility, legitimacy, and complementarity.

- *Credibility.* Individuals often need confirmation from others, especially from trusted sources, before innovations attain credibility (Coleman et al., 1966).
- *Legitimacy.* An innovation is often deemed deviant from existing social norms until there is a critical mass of early adopters to give it legitimacy (Johnson et al., 2006; Rogers, 2003).
- *Complementarity.* Adopting an innovation involves considerations about costs and benefits, and the overall value of the innovation can depend on consumption externalities—the number of prior adopters (M. L. Katz & Shapiro, 1985; Rohlfs, 1974).

In general, the family of complex contagion models is based on the idea that the expected utility of adopting a costly behavior depends on the number of different neighboring adopters (Blume, 1993; Galeotti et al., 2010). That is, the individual-level adoption rate is an increasing function of the number of neighbors who have already adopted. In simple terms: some form of "utility" (or credibility or legitimacy) accumulates as more neighbors adopt, and this utility increases the individual's likelihood to adopt as well. This "accumulation" is what we call social reinforcement. The idea of increasing utility is often combined with the notion of a threshold where a single parameter (θ) divides non-adoption from adoption: once enough "utility" has been accumulated (i.e., enough neighbors of an individual have adopted), then the threshold is crossed and the individual will adopt as well (Granovetter, 1978; Centola & Macy, 2007).[1]

We consider two alternative processes that can serve as a basis for comparison that may allow us to quantify the strength of social reinforcement. The first is *incidental information transfer*. This mechanism assumes that signals from multiple sources have no effect whatsoever. It starts from the assumption "why would a redundant signal add anything"? (i.e., conditional on having received the signal once, what is the benefit of receiving it again from a different source?) After all, humans are good at processing information. For example, learning about a 50%-off sale at the supermarket from one source may be sufficient if behavior is driven by awareness alone (stylized example from DiMaggio & Garip, 2012). Hence, it could be a plausible baseline hypothesis that once an individual has awareness about a product or behavior, they decide whether or not to adopt it then. Under this model, receiving another signal from a different source does not affect the adoption decision.

The second alternative is *independent cascades*. It also assumes adoption depends on information transfer, but it allows for transfer to be probabilistic. Under this independent cascade model, the likelihood of adoption increases with exposure to multiple adopting neighbors because there are more opportunities for the signal to be transmitted. Despite a constant[2] incidence rate, the chance that an individual adopts increases with the number of neighboring adopters simply because the likelihood of successful transmission increases (i.e., the likelihood of observing at least one neighbor's adoption increases if there are more neighboring adopters). This process is assumed in biological contagion, where a constant probability determines the likelihood that a susceptible individual will be infected from contact with one contagious individual (Anderson & May, 1992). Thus, Anne is more likely to get the flu if she interacts with Joe and Nancy who are both infected, than if she interacts with just one of them. Both the *incidental information transfer* and the *independent cascades* model have variously been called simple contagions in previous work (Guilbeault & Centola, 2021). We focus here on the

---

[1] Note, the idea of a threshold depends on increasing utility, but the reverse is not necessarily true. Even in the presence of increasing utility, adoption could still follow some non-threshold pattern even though the empirically observed adoption behavior on the individual-level is binary (adopt vs. not adopt).

[2] When we say "constant", we do not necessarily mean that the infection rate is homogeneous for all individuals in the population (or homogeneous for all ties linking a susceptible individual to an infected one). Our point is that the infection rate is constant with regards to the number of neighboring adopters.

independent cascade (probabilistic) model, since it is the more general one (i.e., circumstances where human behavior is deterministic are quite rare; Wan et al., 2024).

How can we turn these alternative modeling approaches into an empirical measure? A common reading of the early threshold model has given rise to the view that the return from the k^{th} neighbor should be "twice" as large as the $k-1^{\text{th}}$ (Dodds & Watts, 2005, p. 596), that spread in a population has to be "superlinear" (Dodds, 2018, p. 77), or that adoption dynamics should be "non-submodular" (Gao et al., 2019, p. 2) to indicate complex contagion. But in reality, thresholds vary and people may even adopt with below-threshold exposure. Research, including some of the early work by Centola and Macy (2007), has shown that some of the key insights of complex contagion hold even when thresholds are not deterministic. However, research has also pointed to important differences. Wan et al. (2025) show that whereas faster spread on clustered networks is possible, it is not the dominant pattern: even with significant social reinforcement, behavior often spreads faster on random networks. As a result, network-level spreading patterns such as "faster spread on clustered networks" can fail to identify complex contagions even in the presence of strong social reinforcement (in the Appendix, we show that neither "doubling" nor "superexponential" are reliable criteria and can fail to detect positive social reinforcement). Our contribution provides a bridge between earlier conceptual models of complex contagion (Centola & Macy, 2007) and more recent probabilistic formulations (Eckles et al., 2024; Wan et al., 2025) as well as hypergraph models that capture higher-order interactions and individual-level heterogeneity (St-Onge et al., 2021). We propose a unifying empirical criterion to adjudicate among these perspectives by mapping multibody interactions—exposure to multiple infected neighbors—to nonlinear infection rates and contrasting this with the independent cascade (simple contagion) as a baseline. This approach quantifies the complementarities among signals from multiple sources and compares them to the benchmark of independent pairwise interactions (repeat exposure of the same individual).

In the next section, we develop a flexible formulation of complex contagion as a function of social reinforcement, building on insights from recent work (Eckles et al., 2024; Wan et al., 2025) to relax the assumption of deterministic thresholds and instead treat social reinforcement as a continuous effect whose strength can be estimated (St-Onge et al., 2021). Specifically, we formalize complex contagion as a collective mechanism operating over an individual's entire neighborhood, mapping multibody interactions onto nonlinear adoption rates. Importantly, our model preserves the crucial distinction between multiple exposures and exposure to multiple sources highlighted in the original conceptualization (Centola & Macy, 2007). It also integrates evidence for negative reinforcement—where repetitive signals from the same source reduce adoption likelihood (Mønsted et al., 2017)—which classic threshold models do not accommodate. Together, this provides a novel empirical criterion for rigorously quantifying how credibility, legitimacy, and complementarity accumulate through social reinforcement.

An Empirical Criterion for Complex Contagions

The fundamental building blocks of diffusion models are assumptions about how people change their behaviors in response to others who they observe or interact with (Figure 1). Conceptually, without social reinforcement, contact with two *different* neighbors should have the same effect as contacting the same neighbor *twice*. With social reinforcement, on the other hand, the probability of adoption is higher given contact with two different neighboring adopters compared to contact with the same neighbor twice. We use this distinction to develop a criterion for empirical analysis of diffusion processes. Let the variable k denote the number of neighbors who have already adopted a given behavior—we call them "activated" neighbors. Whether or not diffusion will be "complex" vs. "simple" will depend on the relative changes in the observed sample-level adoption between $k = 0, k = 1$ and $k = 2$.

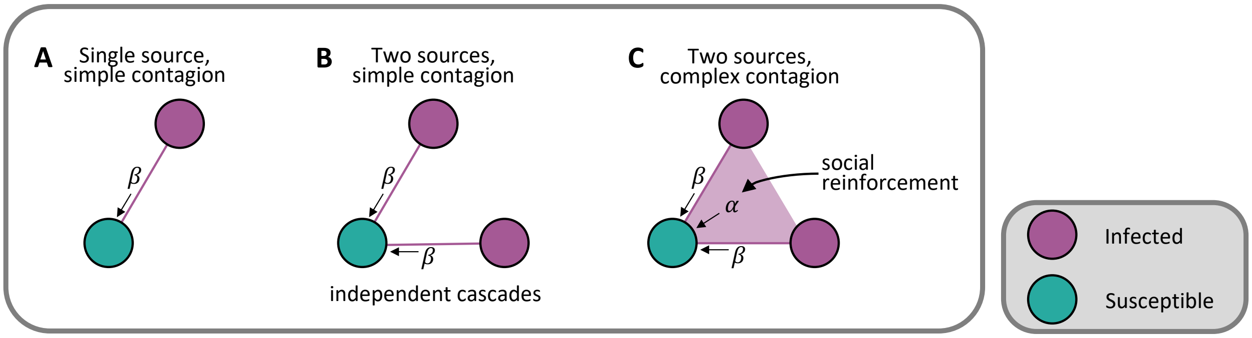

Figure 1. Conceptual distinction between simple contagion processes with one (A) and two (B) infected neighbors and complex contagion with social reinforcement (C).

We formalize this as follows. Every edge between an individual who has already adopted ("infected") and an individual who has not yet adopted ("susceptible") has a probability β of transmitting adoption (infection). In simple contagion models, each edge has an independent (and identical) probability of transmitting adoption (Figure 1A and B). That is, under simple contagion with independent cascades, the probability of adoption with k neighboring adopters is $1 - (1 - \beta)^k$ (i.e., infection is the result of independent pairwise interactions). Building on recent work on hypergraphs (St-Onge et al., 2021), the probability of adoption P_k for a susceptible individual who has k infected neighbors given social reinforcement (Figure 1C) can be described by

$$P_k = \left(1 - \left(1 - \beta^{\frac{1}{\alpha}}\right)^k\right)^\alpha \qquad (1)$$

where α is the social reinforcement parameter.[3] This model flexibly captures different forms of social reinforcement. If $\alpha > 1$ the equation describes a complex contagion process with *positive*

---

[3] We use $\frac{1}{\alpha}$ in the inner term because otherwise α would affect the probability of infection even when $k = 1$. This equation is based on new insights that, at the individual level threshold dynamics combined with heterogeneous temporal patterns lead to infection probability that are approximately power-law (St-Onge et al., 2021). As a result, our equation of the infection probability describes an effective diffusion mechanism that captures a mix of social reinforcement combined with unobserved heterogeneities through a hypergraph structure. While solving this

social reinforcement. If $\alpha = 1$, the equation collapses to the usual simple contagion of independent cascades and describes a process *without* social reinforcement. With $\alpha < 1$ the equation describes *negative* social reinforcement. When α is greater than 1, the contagion process can no longer be seen as independent transmissions but instead "collective" transmissions in which signals from multiple sources reinforce each other. We illustrate this equation visually and show both positive as well as negative social reinforcement (Figure A1). We also show with a simple simulation of a standard diffusion process on a random network that even in the presence of positive social reinforcement, the contagion process in the population overall does not necessarily follow a superexponential regime (Figure A2). Hence, network-level prevalence curves can fail to convey information about positive social reinforcement (this has also been shown by Wan et al., 2025).

Notice that contact with *more* activated neighbors always *increases* the overall adoption likelihood even without social reinforcement ($\alpha = 1$). However, there are diminishing returns for additional contacts, as reflected in the k exponent in the equation. Consider the following example. Assuming some typical $\beta = 0.1$ (10% adoption likelihood), at $k = 2$ this gives us an adoption likelihood for a susceptible individual of $P_{k=2} = 1 - (1 - 0.1)^2 = 0.19$. Note that this is not double the case of single exposure, which would be 0.2. The return of the second exposure is only 0.09 (i.e., $P_{k=2} - P_{k=1} = 0.19 - 0.1 = 0.09$). This can be seen as the *expected return under a simple contagion model* which serves as a counterfactual basis to judge whether an observed adoption likelihood at $k = 2$ shows signs of positive social reinforcement ($\alpha > 1$). Note that while we used $k = \{0, 1, 2\}$ throughout this simplified exposition, our criterion is general and can be applied to compare any P_k to P_{k+1}.

The test of social reinforcement, then, involves testing the sharp null hypothesis of $\mathcal{H}_0\colon \alpha = 1$. To make this directly testable with population-level estimates recovered from a regression model we can translate this into: $\mathcal{H}_0\colon p_{k=2}^{Observed} = p_{k=2}^{Simple} = 1 - (1 - p_{k=1}^{Observed})^2$. That is, we can compare the observed return from contact with $k = 2$ neighboring adopters to the counterfactual expected return under a simple contagion model (with $\alpha = 1$) to determine if a population-level diffusion process is consistent with simple contagion or shows signs of social reinforcement. In practical terms, we estimate β as the population-level average $\hat{p}_{k=1}$ of adoption probability for all individuals with one activated neighbor. We use this estimate to calculate what can be expected at the theoretical adoption probability $p_{k=2}^{Simple}$. We also estimate the population-level average $\hat{p}_{k=2}$ of adoption probability for all individuals with two activated neighbors. Finally, we compare the theoretical adoption probability $p_{k=2}^{Simple}$ to the observed adoption probability $\hat{p}_{k=2}$ to determine whether there is evidence of social reinforcement. We can also construct a Wald-style test statistic that allows us to compute a *p*-value for the test of social

---

equation for α does not have a convenient solution, we can compute a solution numerically substituting values for an observed adoption probability P_k and infection probability β (i.e., $P_{k=1}$) to recover the strength of social reinforcement α.

reinforcement. We formalize this test as a Wald-test style inequality test statistic $W = \frac{(\hat{\theta}-\theta_0)^2}{var(\hat{\theta})}$ and test statistical significance between the estimated $\hat{p}_{k=2}$ and the counterfactual $p_{k=2}^{Simple}$. If the derived p-value is less than, say 0.05, then the null hypothesis is rejected, implying presence of social reinforcement. Otherwise, it means there is not enough evidence to reject $\mathcal{H}_0$. One important implication of our conceptualization is that adoption likelihood at $k = 2$ does not have to be "(at least) double" or "super linear" to indicate signs of complex contagions—which has been used as criterion in past research (Centola & Macy, 2007; Eckles et al., 2024). Any increase above $p_{k=2}^{Simple}$—any $\alpha > 1$—indicates the presence of positive social reinforcement.

By sharpening the "null" of how we can test for the presence of positive social reinforcement on the individual level, we have established an empirically testable complex contagion criterion on the population level.

We do not claim that all behaviors should be subject to complex contagion. Yet we argue that social reinforcement will very likely unfold in the realm of new product adoption—our study context. Consumers typically experience prepurchase uncertainty, which makes them susceptible to social influence (Lee & Bell, 2013; Moe & Trusov, 2011). That is, people facing various risks associated with a new product tend to hesitate until they watch what others do (Conley & Udry, 2010; Duflo & Saez, 2003). The riskier and more costly consumers perceive a new product, the more likely that additional exposure through social contacts will increase the credibility, legitimacy, and complementarity necessary for adopting costly behaviors. Formally:

$\mathcal{H}_0: \alpha = 1.$

H1 [Complex Contagions]: Exposure to two activated peers will exert a higher return on the likelihood of product adoption than the return expected under simple contagions.

Embeddedness

The core claim behind why contagions may be complex is that influence propagated along redundant ties provides valuable social reinforcement (Centola, 2010; Centola & Macy, 2007). Yet it remains unclear as to whether social reinforcement simply requires multiple activated neighbors or if cohesive social structures are necessary to support it. Accordingly, we consider the role of structural embeddedness in conditioning the effects having more than a single activated neighbor for behavior adoption. This consideration is motivated by theoretical and empirical work suggesting that peer effect heterogeneity arises from dyadic characteristics between people and in particular structural embeddedness (Aral & Walker, 2014; Coleman, 1988; Granovetter, 1973; Rajkumar et al., 2022).

Embeddedness points to the social network structure surrounding a dyad: the number of thirds that two individuals share in common (Easley & Kleinberg, 2010).[4] Research in sociology and economics has focused on embeddedness as a relationship characteristic, correlated with yet distinguished from the extent of dyadic interactions, that moderates how two individuals influence each other: the more embedded the relationship of two individuals is, the more they influence each other (Burt, 1987; Granovetter, 1985; Uzzi, 1996, 1997). Having common thirds means that the individuals are nested within a structurally cohesive group with strong solidarity and psychological attachment (Moody & White, 2003). The relationship between two individuals can also be monitored by common thirds. For example, non-cooperative behavior can be quickly identified and sanctioned. Hence, embeddedness is associated with greater levels of trust, which can enhance the perceived credibility of information shared over embedded ties (Uzzi, 1996, 1997). Research continues to reveal that social influence varies depending on the level of embeddedness, for example in the adoption of a Facebook application (Aral & Walker, 2014), finding new jobs (Rajkumar et al., 2022), and online content contribution behavior (Rishika & Ramaprasad, 2019).

Embedded relationships support the confirmation, legitimacy, and credibility needed to adopt costly behaviors. In other words, embeddedness may be the mechanism through which social reinforcement increases adoption. When social reinforcement comes from a structurally embedded network tie, we would expect a stronger effect due to the additional confirmation, legitimacy, and credibility of a product's utility. In contrast, if the second exposure occurs from a non-embedded tie, then the strength of social reinforcement may not be qualitatively different and thus would only be redundant. In this sense, embeddedness can function as the microstructure surrounding a contagion and determine whether an additional exposure to activated peers facilitates adoption due to significant social reinforcement.

H2 [Embeddedness Moderation]: The effects of the number of activated peers on product adoption will be moderated by the number of common thirds.

Methods

Experimental Design and Procedures

We designed a country-scale randomized controlled network experiment in a large Asian country. We partnered with the country's largest cellular phone carrier which had about 50 million subscribers. The experiment was part of a viral marketing campaign for a mobile data product. The campaign offered a sharable coupon code that offered one-time usage of 60 megabytes of free mobile data. The coupon contained a unique code for activation through a text

---

[4] Embeddedness in this sense is a dyad-level construct. It is an edge-level measure sometimes called edge-embeddedness. Clustering coefficient similarly captures the node-level degree of clustering (Watts & Strogatz, 1998).

message.[5] After seeding coupon codes in the social network of customers, we observe product adoption for one week. At the time of the experiment 97% of customer connections were prepaid and 98% were 2G. We chose this setting to study viral word-of-mouth marketing in a realistic context. In this setting we cannot directly induce customers to share our viral message—we can only encourage them to share. That is, our experimental design follows an encouragement design (Bradlow, 1998): we encourage a random set of so-called seed customers (or seed nodes on the network) to adopt and share our coupon code and then study the spillover effects of product adoption among their neighbors who are the focal nodes of our study. Encouragement designs are widely used when researchers are interested in the effect of behaviors that are not under direct control of the experimenters (Eckles et al., 2016). Our experiment is also a form of "inside out" experiment (Aral & Walker, 2011) where the focal observations of our analysis are not the encouraged customers themselves, but rather their neighbors. Such a design is both unbiased and efficient (Aral & Walker, 2011), as these focal customers are randomly exposed to treated neighbors who have been encouraged to share a coupon code.

To track spillover effects, we constructed the entire social network of all customers (i.e., all active SIM cards) by using detailed records of call volumes (frequency and length) and text message frequencies over a three-month period. That is, our social network includes information about tie strength based on contact frequency (SMS and call frequency) and contact duration (voice minutes) of all customers. The network consists of 46M nodes and 567M weighted edges. A social network constructed from detailed call records is ideal for this study for two reasons. First, prior research has demonstrated high accuracy between self-reported friendships and proximity, location, and contact time derived from call records (Eagle et al., 2009). Second, SMS text messages were the primary medium for the product's adoption and transmission. The coupon code was simultaneously delivered to all seed customers via SMS messages, and users could redeem the codes by sending the code as text messages to a designated number. SMS text messages were also the most likely medium to share coupon codes with friends as it consisted of random alphanumeric sequences that would have been hard to transmit verbally without error.

---

[5] Coupon codes could be shared by anyone who received them (seeded or not) including to neighbors who are multiple hops away, but our analysis is focused on those focal nodes who are directly connected to seed nodes. Sharing a coupon code is each individual's own decision—whether to share at all or who to share with—and there was no automated broadcast. Each code could be adopted an unlimited number of times but every customer was only allowed to redeem one code from this campaign.

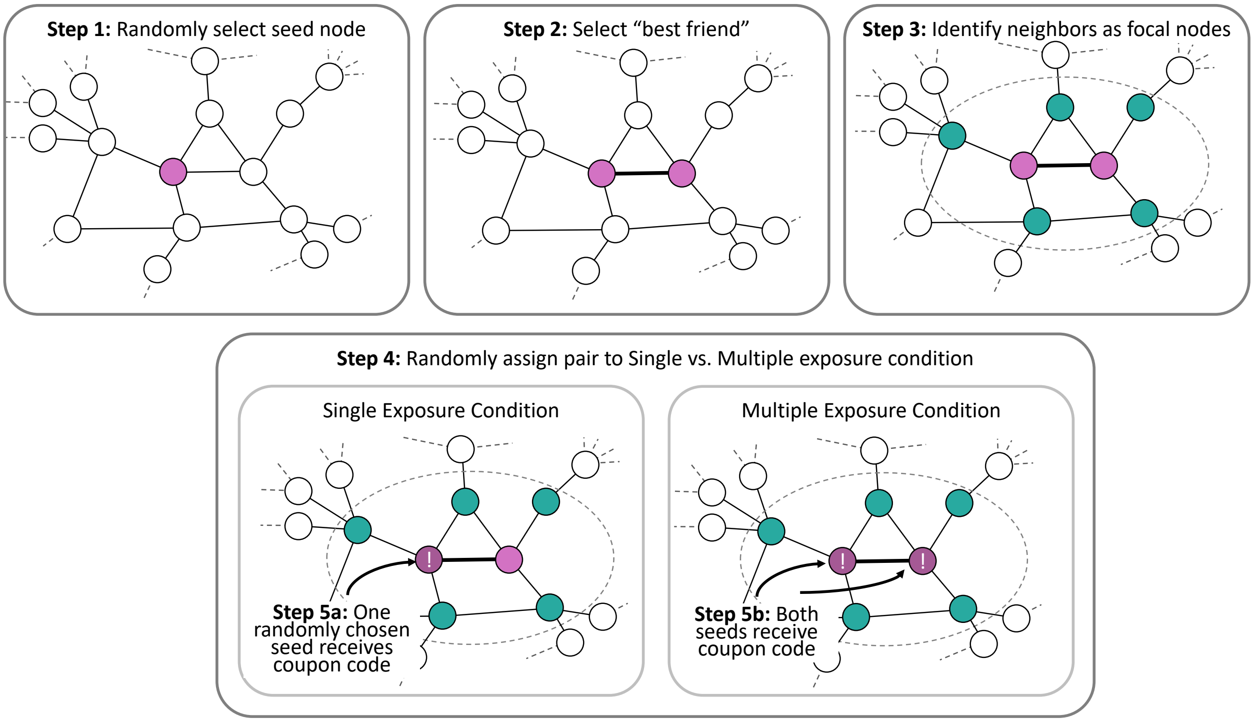

Figure 2. Experimental design.

Within the general framework of this peer encouragement experiment, we randomly sampled pairs of two connected customers and randomly assigned them to two treatment conditions (Figure 2): one in which one randomly chosen customer of the pair received a sharable coupon code and one in which both customers of the pair received a sharable coupon code. This experimental design creates random variation among focal nodes (blue nodes in Figure 2) with regard to how many treated neighbors they have who received a coupon code which allows us to test complex contagion. The list below illustrates the steps involved in the experimental design in more detail.

1. We sampled a set of random nodes among eligible[6] target customers as our seed nodes.
2. Among the randomly selected seed node's immediate neighbors who are also in the set of eligible target customers, we picked the "best friend", based on the highest amount of prior conversation minutes edge weight in the social network.
3. Identify all immediate neighbors of seed nodes as "focal nodes" which form the observations of our analyses.
4. Randomly assign seed pairs to the "single exposure" or "multiple exposure" treatment condition.

---

[6] Eligible criteria included customers who owned an Internet-enabled device but who were not already subscribing to a data plan package and were minimally active users as indicated by having at least four incoming and four outgoing SMS messages in the 3-month period before the experiment. Targeting this group was necessary for practical purposes to actually observe the processes of product diffusion through a network.

a. In the single exposure condition, we randomly selected one of the seed pair customers who would receive a sharable code.
b. In the multiple exposure condition, both seed customers received a sharable code.

This experimental design was deliberately chosen and designed to maximize efficiency and reduce bias. It has several key advantages over other alternative designs. First, the inside-out design is efficient since all neighbors of treated nodes become observations in the experiment. Second, seeding codes in pairs generates more observations of focal nodes with two treated neighbors than seeding codes completely randomly in the population would. That is, ours is an efficient design that increases the statistical power over alternative random seeding strategies that do not use pairs of connected nodes. Third, selecting seed pairs based on "best friend" relationship avoids known bias from over-sampling of high-degree nodes that selecting connected neighbors randomly would (due to friendship paradox; Jackson, 2019). Overall, this design creates efficient, unbiased random variation in exposure to one versus two treated neighbors. The focal nodes connected to the treated and the untreated seed are ex ante equivalent—including their embeddedness. Overall, this experiment generated observations for about 4 million focal nodes. Overall adoption was 12.8% among seeds and 0.8% among focal nodes. We summarize our key variables and measures in Table A1 and provide descriptive statistics in Table A2. We discuss additional challenges we encountered while conducting this experiment in Section 3 of the Appendix.

Empirical Model

Despite their central importance to social science, causal peer effects are notoriously difficult to estimate reliably in empirical data (Aral, 2011; Manski, 2000; Shalizi & Thomas, 2011). Peer effects are often confounded with homophily and common external causes, inference is challenging in network settings where treatment units are not independent, and treatment interventions designed to elicit peer effects do generally not deterministically spread to these neighbors (compliance). The analysis is further complicated as different stakeholders such as marketing managers judging the practical efficacy of a novel marketing campaign and social scientists interested in testing specific hypotheses of causal effects may be interested in different quantities. As a result, even in the presence of a carefully designed randomized experiment, no single analysis captures all relevant aspects. We provide comprehensive evidence from three different, but complementary, analyses. Our first analysis (intent-to-treat) helps assess the overall efficacy of the marketing campaign overall, our second (encouragement design analysis using instrumental variables) addresses compliance, and our third (exact p-value approach) addresses non-independent observations in network settings.

Analysis 1: Intent-to-Treat Effects

Our first set of analyses estimates intention-to-treat effects (Gerber & Green, 2012). Here we estimate the likelihood that focal node i adopts the product based on the randomly assigned treatment status of its connected seed nodes: having zero, one, or two treated neighbors. This analysis directly exploits on the random variation generated by our experimental design. A key strength of this analysis is that it allows us to estimate causal differences in the net impact of the intervention and that it is straightforward to interpret. Furthermore, this analysis allows us to look at product adoption within a narrow time window after seeds received their coupon code SMS text messages. This allows us to investigate what role—if any—indirect spreading in the network via other indirect paths might play (choosing a short time window reduces the likelihood that codes spread across indirect paths between the seed node and the focal node). However, one weakness of this analysis is that it leaves unanswered some questions about the mechanism of the observed effects. Specifically, since seed nodes were only encouraged to share coupon codes with their neighbors, but actual spreading is unobservable in this field setting, our second analysis analyzes our data as an encouragement design.

Analysis 2: Encouragement Design with 2SLS Instrumental Variables

With the increased interest in peer effects but the challenge of causal identification in observational data, recent work increasingly focuses on randomized experiments to estimate how the behavior of an individual is affected by the behavior of their peers (Bakshy et al., 2012; Moffitt, 2001; Walker & Muchnik, 2014). These experiments often relied on randomly manipulating the *mechanism* by which the peer behavior is transmitted to the focal individual (Eckles et al., 2016). For example, Aral & Walker (2011) use a custom Facebook app to randomize which peers are sent viral messages about their friends' behavior (adopt a product). Experimental designs that manipulate the mechanism that deterministically sends messages from individuals to their peers are often not possible, especially in natural word-of-mouth marketing settings. In general, directly randomizing the behavior of existing peers in field experiments is not possible or desirable (Eckles et al., 2016). Instead, we can think of randomized designs that encourage individuals to share information with their friends in hopes of influencing their behavior. Our experimental design can be thought of as such an encouragement design. It randomly assigns the peers of focal individuals to a treatment in which they are encouraged to adopt and share a mobile internet product. If those seed individuals comply with the encouragement, we may then observe potential spillover effects on their peers (the focal individuals of our study). Such encouragement designs are often used in social science when behaviors cannot be directly controlled. For example, in political science, one may encourage voters to watch a debate on TV (Albertson & Lawrence, 2009). Even though not all individuals may comply with the intervention of watching the debate, we can then still estimate the causal effect of the intervention on voting behavior for those who have randomly received the encouragement (i.e., for compliers). Such an encouragement design can be adapted to study peer effects in product diffusion on networks (Eckles et al., 2016) by randomly assigning the *peers* of focal individuals to an encouragement to adopt and share that product with their friends.

That is, seed individuals are randomly assigned to an encouragement Z_i. Then the endogenous seed behavior K_i is observed as well as the behavior of the focal individuals Y_i. We can then examine how the encouragement to adopt mobile internet "spills over" to the behaviors of focal individuals. In our setting, we do not observe communication content and hence we do not observe "sharing" of coupon codes directly. We do observe whether seed nodes redeem the coupon codes and we take this “activation” behavior as the source of the spillover effect.[7] Throughout, it is important to keep in mind that the sample of this analysis are focal individuals: those nodes in the network who are connected to seed nodes who were randomly assigned to an encouragement condition. Consequently, i always refers to focal individuals, while K_i summarizes the endogenous activation status of the seed nodes that focal individual i is connected to. We estimate the effect of having K_i activated neighbors on Y_i by using the random variation in K_i caused by assignment to the peer encouragement, Z_i. In terms of the instrumental variable estimation, this means we have two binary endogenous variables: $K_i = 1$ and $K_i = 2$.

This potential outcome formulation requires four assumptions (Eckles et al., 2016): (1) an exclusion restriction for instrumental variables (IVs); (2) no interference; (3) direct-effect-bounded interference; and (4) monotonicity, which is that there are no defiers. The exclusion restriction seems plausible in our case as participants are blinded to their assignment. It seems safe to assume that participants did not realize they were randomly assigned to treatments—especially given the minimal nature of the encouragement—or if they were it was just a normal kind of marketing campaign that would not otherwise affect their behavior (see more in depth discussion steps to mitigate remaining concerns in Section 5 of the Appendix).

To estimate the causal effect of K_i we estimate the following regression equations using two-stage least squares (2SLS):

$$1(K_i = 1) = \gamma_1 {Z_i}^R + \gamma_2 X_i + \eta_i \tag{2a}$$

$$1(K_i = 2) = \gamma_1 {Z_i}^R + \gamma_2 X_i + \eta_i \tag{2b}$$

$$Y_i = \beta_1 1\big(\widehat{K_i = 1}\big) + \beta_2 1\big(\widehat{K_i = 2}\big) + \gamma X_i + \epsilon_i \tag{3}$$

where we instrument for the two binary endogenous variables $K_i = 1$ and $K_i = 2$ using the randomly assigned encouragement (Z_i^R) of i’s connected peers. We use four binary instruments that stem from the combination of our random assignment and status of the connection to seed pairs: (i) single exposure & tie with one seed; (ii) single exposure & tie with two seeds; (iii) multiple exposure & tie with one seed; and (iv) multiple exposure & tie with two seeds. The first-stage model predicts the propensity of i’s peers to adopt the product, estimating the propensity of having one activated peer $1(K_i = 1)$ and the propensity of having two activated peers $1(K_i = 2)$. The second-stage model predicts the focal node’s adoption based on the social influence of having one or two activated neighbors. X_i is a vector of control variables.

This estimator is consistent and unbiased when Z_i is randomly assigned and additional covariates are included in the 2SLS model (Aronow & Carnegie, 2013). Intuitively, 2SLS

---

[7] As many peer effect studies before us, we study spillover effects of behavior: how does the behavior of one individual (adopt) affect the behavior of another individual (also adopt). We theorize that “information transfer” is the most plausible mechanism through which this peer effect materializes. However, since we do not observe communication content we cannot verify information transfer as the mechanism underlying the peer effect in our study. Some other plausible contributing mechanisms include changes in preferences or increased utility (if more peers are using mobile internet then ego has more communication partners). For the purpose of our paper it seems plausible to think of “activation” as a proxy measure for “sharing”.

rescales the intent-to-treat (ITT) effect of Z_i on Y_i (the average effect of treatment assignment on the outcome) while acknowledging that only compliers will respond to changes in Z_i. This estimator is the average causal treatment effect for compliers (LATE).

The LATE estimates reflect two-sided non-compliance: (a) some seeded individuals did not adopt despite being encouraged (i.e., never-takers); and (b) some seeded individuals may adopt the product despite not being treated directly (i.e., always-takers). Since adoption overall is quite low there cannot be many always-takers. This offers good partial identification, and we are hence mostly concerned with potential bias from never-takers: encouraged seeds who did not adopt. It could be that we observe social reinforcement from two activated peers (*k*=2 cases) predominantly in those regions of the social network which are especially interested in the mobile internet product and thus among individuals who are especially susceptible to adopt. This unobserved heterogeneity in product interest may be a confounder in our analysis making LATE overestimate the effect of social reinforcement given a positive correlation between exposure to two activated peers and the focal node adopting.

We address this by estimating the (population) Average Treatment Effect (ATE), instead of the LATE, with the reweighting method developed by (Aronow & Carnegie, 2013) and the software implementation provided in (Aronow et al., 2018). The intuition behind the method is analogous to inverse probability weighting (IPW) for sample correction. Specifically, the method leverages the random assignment of the instrument to compute a consistent estimator of the ATE for compliers and then reweights the population so that the compliers have a covariate distribution that matches that of the entire population. One complication of this approach in the current setting is the definition of complier strata.[8] In the canonical case of a single, binary endogenous variable, there is a single latent stratum of compliers; but here are multiple such strata, since focal nodes have two peers. We define a single complier group for focal nodes who, if both their peers are assigned to treatment, will have two adopter peers (we call these "full compliers"). We calculate a maximum likelihood estimate of this full complier strata for each focal node, given the random treatment assignment (receiving a coupon code or not) and covariates (degree and local network density). We then reweight the sample in the 2SLS analysis using the inverse compliance score weighting (ICSW).

Analysis 3: Inference in Network Experiments

The ITT and 2SLS IV approaches outlined above are tailor-made for the analysis of encouragement designs. However, both approaches rely on the assumption of no interference (SUTVA), a well-documented problem for networked experiments (Aral, 2016; Aral & Walker, 2014). The reason behind why ignoring SUTVA biases treatment effect estimates is that such interference means that treatment units cannot serve as counterfactuals to each other, since each unit would have potential outcomes defined differently. Since this assumption is likely violated

---

[8] We want to thank an anonymous reviewer on a previous submission for helpful comments on the issue of complier strata.

in our setting despite being a randomized experiment, we complement those two analyses with a third analysis that specifically addresses the concern that observations in network settings may not be independent. Specifically, we apply the computational method developed by (Athey et al., 2018) to test specific sharp null hypotheses concerning spillover effects under network interference using exact randomization-based p-values (Fisher, 1925). The method requires four ingredients: a mathematically formulated hypothesis, an artificial experiment, a test statistic, and a randomization-inference step. We explain each step in detail below. Throughout, we draw on the following input data: a set of N individuals i (nodes), an outcome vector Y_i, a vector of random treatment assignments W_i, and a network adjacency matrix G with $G_{ij} = 1$ if individual i and j are connected and 0 otherwise.

Hypotheses. A set of mathematically formulated hypotheses allows testing for spillovers, and specific forms that spillover effects might take. We formulate three specific hypotheses (see Appendix for details and formal mathematical formulations): *Hypothesis 0* (No Spillovers), *Hypothesis 1* (No Threshold Peer Effects), and *Hypothesis 2* (No Peer Effect Heterogeneity).

Artificial Experiment. We select half of the nodes at random and designate them as the "fixed" subset P_F, the other half as "variable" subset P_V.

Test Statistic. The test statistic captures the degree to which the unexplained variation (residual) of the null model, for example, a model that allows for a direct treatment effect but no spillover effect, covaries with a term capturing the alternative hypothesis (see Section 6 of the Appendix for equation).

Randomization-Inference. During the randomization-inference step we randomly permute the treatment status of "variable" nodes in P_V, holding the treatment status of "fixed" nodes in P_F invariant. Next, we calculate the number of treated neighbors that the "fixed" nodes have under the permuted treatment status of the "variable" nodes. We then calculate our test statistic for all nodes in P_F. Notice how this "artificial experiment" will thus keep the treatment status of the "fixed" nodes invariant, but randomly reshuffle the treatment status of their neighbors thus generating random variation in the peer effect term. This step is repeated 100,000 times to generate a null distribution of test statistics. Finally, we calculate the exact p-value for the hypothesis as the degree to which the observed test statistic exceeds the null distribution.

Results

Result 1: Intent-to-Treat Effects

Our first set of analyses focuses on spillovers from changes in peer treatment status alone. That is, we start with an analysis of the average effect of treatment assignment on the outcome (ITT) ignoring the more nuanced analysis of the encouragement design for the moment. Since the treatment status (targeting seed consumers) is something that can be controlled by a marketer analyzing the ITT effect can help us determine the efficacy of clustered targeting for new product adoption. Using OLS, we regress product adoption on having zero, one, or two peers randomly receive a coupon code treatment (Eckles et al., 2016). To investigate the effect of indirect spreading across the social network, we use different time windows to determine focal individuals' outcomes of product adoption. For short time windows (6h, 12h, and 24h; Models 1, 2, 3) we find no signs of indirect spreading and adoption likelihood is zero for individuals without immediately treated neighbors (model intercept). We find significant spillover effects from having one treated neighbor. This provides a baseline finding that the viral marketing campaign generated positive peer-to-peer spillovers.

Are there signs of social reinforcement among those who are connected to two treated neighbors? We find much larger significant effects for individuals who are connected with two treated neighbors. We use the regression coefficients of $k_{ITT} = 1$ and $k_{ITT} = 2$ from Model 1 to calculate the adoption probabilities $P_{k=2}$ and $P_{k=1}$. Then, we use a numerical solver (such as WolframAlpha or a custom implementation of a binary search; see pseudo-code in Appendix 10) to seek a solution for the social influence parameter α in equation in Eq. (1). That is we solve $P_{k=2} = \left(1 - \left(1 - P_{k=1}^{\frac{1}{\alpha}}\right)^2\right)^\alpha$ after substituting 0.0008 for $P_{k=2}$ and 0.0001 for $P_{k=1}$. We find $\alpha \approx 3.094$ indicating strong positive social reinforcement. We can also test for evidence of complex contagion against $H_0\colon p_{k=2}^{Observed} = p_{k=2}^{Simple}$using our Wald-style test introduced above. The counterfactual adoption likelihood when exposed to two activated peers under simple contagion is

$$p_{k=1}^{obs} = \beta_{k=0}^{obs} + \beta_{k=1}^{obs} = -0.0001 + 0.0002 = 0.0001$$
$$p_{k=2}^{simple} = 1 - (1 - 0.0001)^2 = 0.0002.$$

The observed adoption likelihood on the other hand is $p_{k=2}^{obs} = \beta_{k=0} + \beta_{k=2} = 0.0008$. We test the statistical significance of the inequality $p_{k=2}^{obs} = p_{k=2}^{simple}$ using a Wald test which yields $W = 4.31$ and a p-value $= 0.038$.

Dependent Variable:	Adopted Product within first ...					
	6h	12h	24h	36h	48h	7 days
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	-0.0001^{***}	0.0000	0.0000	0.0002^{***}	0.0002^{***}	0.0008^{***}
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)
k_ITT = 1	0.0002^{***}	0.0002^{***}	0.0003^{***}	0.0004^{***}	0.0004^{***}	0.0007^{***}
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)
k_ITT = 2	0.0008^{**}	0.0012^{***}	0.0012^{***}	0.0017^{***}	0.0019^{***}	0.0039^{***}
	(0.0003)	(0.0003)	(0.0003)	(0.0004)	(0.0004)	(0.0006)
Controls						
Degree	-0.0000^{**}	-0.0000^{*}	-0.0000^{**}	-0.0000^{**}	-0.0000^{***}	-0.0001^{***}
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Duration	-0.0005^{***}	-0.0007^{***}	-0.0010^{***}	-0.0015^{***}	-0.0017^{***}	-0.0035^{***}
	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Frequency	0.0005^{***}	0.0005^{***}	0.0006^{***}	0.0008^{***}	0.0009^{***}	0.0015^{***}
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Embeddedness	0.0122^{***}	0.0169^{***}	0.0205^{***}	0.0278^{***}	0.0296^{***}	0.0512^{***}
	(0.0009)	(0.0011)	(0.0012)	(0.0014)	(0.0014)	(0.0019)
α	3.0939	2.2488	1.8522	1.7091	1.7023	1.6732
Wald $P^{Obs}_{k=2} = P^{Simple}_{k=2}$	4.3112	4.0913	2.5398	2.8922	3.2804	7.6304
R^2	0.0004	0.0004	0.0005	0.0006	0.0006	0.0009
Adj. R^2	0.0004	0.0004	0.0005	0.0006	0.0006	0.0009
Num. obs.	2,502,187					

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 1. Coefficient estimates of the intention-to-treat effect on focal nodes' product adoption likelihood based on exposure to treated neighbors using different time windows.

Not surprisingly, the rate of indirect spreading increases over time. We find significant indirect spreading and even focal nodes who are not directly connected to a treated neighbor have a substantial likelihood to adopt the product. The conclusion about strong social reinforcement holds; for example, using coefficients from Model 6 we find a similar $\alpha \approx 1.67$ and a significant Wald test ($W = 7.63; p = 0.006$). This finding furthermore remains robust when we iteratively trim observations until all $\hat{y}_i$ lie within $[0,1]$ to address the issue of unbiasedness and consistency in the linear probability model (Horrace & Oaxaca, 2006, see Table A7 in Appendix 9).

Treatment effect heterogeneity. To investigate treatment effect heterogeneity, we introduce an interaction term between the number of treated neighbors and embeddedness (Table A3). Using both the 6h (Model 1) and 7 day adoption window (Model 2) we find significant positive interaction for both the $k = 1 \times Embeddedness$ and the $k = 2 \times Embeddedness$ coefficients. In the Appendix we also show that spreading along indirect paths cannot explain these results. This suggests that structural embeddedness between individuals is a key driver behind social reinforcement (H2).

Result 2: Peer Effects in Networks with Peer Encouragement Designs

We now shift to the analysis of our data as a peer encouragement design using 2SLS. In this analysis, we focus on the degree to which focal individuals are exposed to "activated" peers, which is instrumented by our random treatment assignment (the encouragement). We first examine the effect of the peer encouragement treatment on the exposure of focal nodes to activated neighbors who adopted mobile internet (i.e., first-stage effects). That is, only when the seed node acts upon receiving the marketing message are focal individuals exposed to an "activated" neighbor. All treatments have large and significant direct effects on seed nodes' adoption (i.e., large F-statistics; Table A4). This indicates that our random treatment is a good instrument for degree of exposure to activated neighbors among focal individuals. This randomly induced variation in exposure to activated peers allows us to estimate causal peer effects on the focal nodes' product adoption behavior. For example, having a peer who was assigned to receive the coupon code treatment increases the focal node's number of activated peers by 0.064 (Model 1). These effects are robust to the inclusion of control variables (Model 2). Notably, we find that degree of seed nodes has a significant negative coefficient. This suggests some general implications for diffusion as it may reflect that high degree individuals have bandwidth constraints that reduce the likelihood of transmission to any given alter. This indicates that some important anti-correlation between seeds that are likely to adopt and seeds that can spread to many neighbors.

Second stage results are in Table 2. We present both OLS regression results and 2SLS and results both with and without additional control variables. The table is organized such that with each column we further control for potential omitted variables so that we can learn about the sources and size of any bias. In Model (1), we show baseline OLS results. These are observational estimates of peer effects that decision-makers could rely on in the absence of our peer encouragement design. We find substantial evidence for social spillover: being connected to an activated peer substantially increases adoption likelihood. Being friends with one activated peer increases adoption likelihood by 0.8%, while being friends with two activated peers increases adoption likelihood by 3.3%. We again apply our test of social reinforcement which yields $\alpha \approx 2.0$ ($W = 10.76$; $p\text{-value} = 0.001$). We hence reject the null hypothesis of simple contagion.

Dependent Variable:	Adopted Product				
	OLS	LATE (2SLS)	ATE (ICSW)	ATE (ICSW)	OLS
	(1)	(2)	(3)	(4)	(5)
Intercept	0.002^{***}	0.001^{***}	0.002^{***}	0.001^{***}	0.001^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
k = 1	0.008^{***}	0.010^{***}	0.010^{***}	0.010^{***}	0.008^{***}
	(0.000)	(0.001)	(0.001)	(0.001)	(0.000)
k = 2	0.033^{***}	0.095^{***}	0.099^{***}	0.064^{***}	0.031^{***}
	(0.005)	(0.012)	(0.012)	(0.012)	(0.005)
<u>Controls</u>					
Degree				0.000	-0.000^{*}
				(0.000)	(0.000)
Duration				-0.003^{***}	-0.003^{***}
				(0.000)	(0.000)
Frequency				0.001^{***}	0.001^{***}
				(0.000)	(0.000)
Embeddedness				0.047^{***}	0.048^{***}
				(0.002)	(0.002)
α	2.004	4.268	4.462	3.261	2.064
Wald $P^{Obs}_{k=2} = P^{Simple}_{k=2}$	10.755	40.533	43.061	13.737	10.083
Adj. R^2	0.002	0.001	0.002	0.003	0.003
Num. obs.			2,502,187		

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 2. Regression results of encouragement design. OLS, 2SLS, and ICSW estimates for the effect of activated neighbors on product adoption.

Model (2) exploits random variation in contact with activated peers through our randomized encouragement design. This estimate represents the local average treatment effect (LATE) for focal nodes connected to compliers: seeds who were treated with the encouragement and adopted. That is, it identifies average effects for individuals on the margin where being friends with treated others who have adopted makes a difference. We find that IV estimates from the peer encouragement design are larger than observational estimates (OLS). This suggests that unobserved features, like interest in the product, local network structure, or communication patterns, lead OLS to understate the benefits of redundant exposure. Our test of social reinforcement yields $\alpha \approx 4.27$ ($W = 40.53$; p-value < 0.001) and we hence reject the null hypothesis of simple contagion.

Model (3) shows the results for the (population) Average Treatment Effect using the ICSW method. We find very similar coefficients than our LATE estimates which again indicate the presence of social reinforcement ($\alpha \approx 4.46$; $W = 43.06$; p-value < 0.001).[9]

---

[9] We ran several robustness tests using different definitions of complier strata and different ways of estimating complier scores, all yield consistent results. Specifically, we ran robustness tests using (a) no controls, (b) using local network density as control variable and (c) alternatively estimating complier scores on the level of individual seed nodes and then aggregating it to the level of the focal node by mean or max pooling. Empirically, different ways of defining complier strata and estimating them, result in very similar complier scores (with a correlation of around $\rho = 0.95$) potentially indicating strong homophily among the small network clusters surrounding focal nodes. We also apply the standard Winsorization of very small compliance scores recommended in Aronow & Carnegie (2013).

Model (4) is fully specified with additional control variables for (a) local network structure and (b) tie strength. The spillover effects are somewhat smaller but still large and well within the range rejecting the null hypothesis of no social reinforcement ($\alpha \approx 3.26$; $W = 13.74$; p-value < 0.001). Finally, Model (5) shows another comparison with OLS estimates which are very similar to the estimates without controls. As robustness tests, we report ATE estimates with the varied time windows as in our ITT analysis in Table A5 in the Appendix.

We summarize our key test of Hypothesis 1 (Complex Contagion) visually in Figure 3. Our main focus lies in the endogenous conditions that seeds' adoption behavior creates for focal nodes: how much does exposure to two activated neighbors increase an individual's adoption likelihood, as compared to exposure to just one activated neighbor? And is this increase higher than the counterfactual increase we would expect of exposure to two activated neighbors under a counterfactual simple contagion model? For both the estimates of observational effects (OLS) and the peer encouragement design (2SLS), we find that exposure to two activated neighbors is significantly higher than the estimate under a counterfactual simple contagion. This finding supports the hypothesis of positive reinforcement and provides a strong indication of complex contagion.

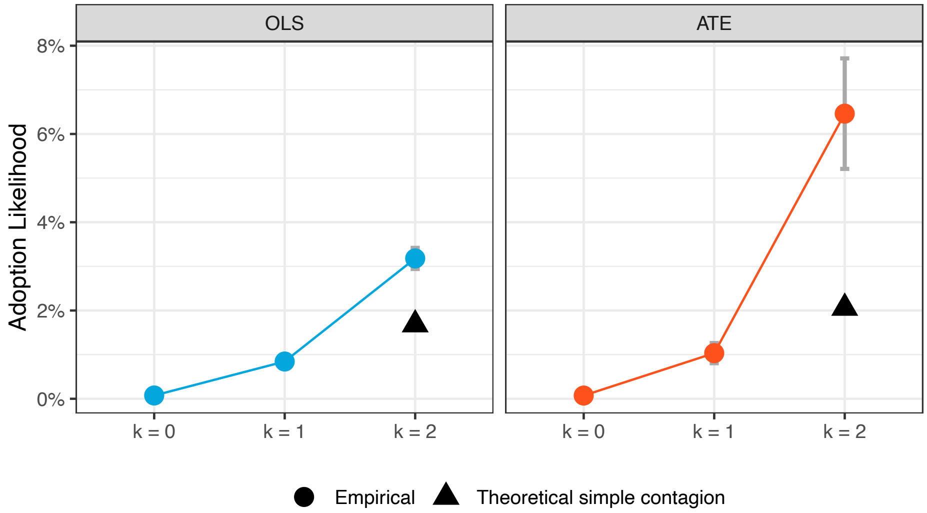

Figure 3. Comparison of estimates using observational analysis using OLS (left) and ATE (right). Coefficients are from Model (5) and Model (4), respectively from Table 2.

How does the size of social reinforcement estimate vary across time? We plot our social reinforcement measure α using different time windows (Figure A5) and find that it becomes larger over the course of the viral marketing campaign. This suggests that "laggards," as compared to "early adopters," of Roger's adopter categories (Rogers, 2003) may be more susceptible to redundant exposure through social contacts.

Heterogeneous Treatment Effects: The role of embeddedness

Is there any treatment effect heterogeneity with regard to the number of common neighbors that a focal individual has with those activated peer—their structural embeddedness? We extend our main analysis by including the interaction term of the number of activated seeds and embeddedness (Table 3). The models are similarly organized as our main 2SLS analysis, estimating both LATE and ATE effects. Yet these models differ in that we now enter k as a single count variable instead of three separate dummies for having no $(k = 0)$, one $(k = 1)$ or two $(k = 2)$ activated peers. This choice is made by our change in focus from testing complex contagion hypothesis to analyzing effect heterogeneity. Furthermore, interacting the embeddedness variable with all the discrete terms of k would cause us to run out of instruments in a 2SLS IV model.

Dependent Variable:	**Adopted Product**				
	LATE (2SLS)			ATE (ICSW)	
	(1)	(2)	(3)	(4)	(5)
Intercept	0.001***	0.000*	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
k (continuous)	0.018***	0.014***	0.004	0.002	0.002
	(0.001)	(0.001)	(0.003)	(0.003)	(0.003)
Embeddedness		0.046***	0.018	0.011	0.011
		(0.002)	(0.010)	(0.011)	(0.011)
k × Embeddedness			0.305**	0.368**	0.366**
			(0.107)	(0.112)	(0.112)
Controls					
Degree					0.000
					(0.000)
Duration					−0.003***
					(0.000)
Frequency					0.001***
					(0.000)
R^2	−0.001	0.002	0.005	0.005	0.005
Adj. R^2	−0.001	0.002	0.005	0.005	0.005
Num. obs.			2,502,187		

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 3. 2SLS and ICSW regression of effect heterogeneity.

2SLS Models (1) and (2) show that embeddedness as a main effect has a small but significant positive effect, with a slight decrease of the k effect. Model (3) is our main model of effect heterogeneity and adds an interaction term between the number of activated peers an individual has and the structural embeddedness with those seed individuals. We find a very strong effect of the interaction effect $(\beta = 0.305; p < 0.01)$ while neither the main effect of k nor the main effect of embeddedness are significant. Model (4) shows estimates of the ATE, which are very similar, indicating that the LATE model controls well for heterogeneity in the effect. Including additional controls in Model (5) barely changes the main coefficient of interest. Overall, we find (descriptive yet not causal) support for our second hypothesis (H2) that the effect of activated peers on product adoption depends on support from local network structure: embeddedness is crucial for spillover effects to materialize.

Results 3: Exact P-Values for Network Interference

To complement the above regression results, we test several sharp null hypotheses through exact p-values. First, we test the null hypothesis of no spillover using the score statistic from Eq. A1 in the Appendix 6. The resulting p-value is small: 0.006 (Figure A5, left) leading us to reject the null hypothesis of no spillovers. Second, we test a baseline of our Hypothesis 1 of social reinforcement through a hypothesis of threshold peer effects (Eq. A2). Intuitively, the score statistic tests the null hypothesis that there is no effect of the number of treated neighbors: it only matters whether an individual has any treated neighbors. We find a p-value = 0.018 (Figure A5, middle) leading us to accept the alternative hypothesis that the number of treated neighbors matters. This is a weaker hypothesis than our hypothesis of social reinforcement (H1) since it does not explicitly test for the difference in effect size of having one versus two treated neighbors. However, it serves as an important baseline test using a method that explicitly accounts for the connected nature of observations: clearly, we can only find support for social reinforcement if the number of treated neighbors matters. The key question answered by the two analyses above was: How much larger is the effect of exposure to two neighbors? For which we find in both that the effect is much larger than would be expected under simple contagion. Third, we test our Hypothesis 2 of the role of embeddedness. We test the null hypotheses of no peer effect heterogeneity: we test that only the number of treated peers matters, not which peers are treated (Eq. A3). Specifically, we test that the embeddedness—number of common friends—of the tie between the focal individual and the treated neighbor does not matter. We again find a small p-value of 0.016 (Figure A5, right). This leads us to accept the alternative that it matters which neighbor is exposed to the treatment supporting Hypothesis 2.

Discussion and Contribution

How behavior spreads from person to person is a central question in social science (Hinz et al., 2011; Gelper et al., 2021; Godes et al., 2005; Godes & Mayzlin, 2009; Kumar & Sudhir, 2021; Tellis et al., 2019; Goldenberg et al., 2001; E. Katz & Lazarsfeld, 1955; Manski, 2000; Mayzlin, 2006; Rogers, 2003; Young, 2011). The central question guiding our study is whether behavioral adoption increases through social reinforcement. We find robust causal evidence for positive social reinforcement across all three of our analyses offering strong support for the presence of complex contagion: the contagion process cannot be understood as independent cascades, but rather as a process in which signals from multiple sources amplify each other through synergistic interdependence (H1: Complex contagions). We consistently find evidence of positive social reinforcement ($\alpha > 1$) in both the ITT and encouragement design analyses as well as support for a threshold hypothesis using the exact p-value method. This provides the most comprehensive evidence for the role of social reinforcement in diffusion to date. We also have some evidence of anti-correlation between seeds that are likely to adopt and seeds that can spread to many neighbors, and that the strength of social reinforcement is larger for late adopters. Across our analyses we also find that these treatment effects are heterogeneous with regard to more structurally embedded neighbors (H2: Embeddedness). We acknowledge, however, that this embeddedness finding remains suggestive since embeddedness was not experimentally manipulated (which is implausible in any network experiment relying on existing ties).

Why did product adoption show signs of social reinforcement? The context of our study was the adoption of a new product—the use of mobile Internet—for customers who were not already accessing the Internet this way. Even though adoption of the 30-day trial was free, customers may have been hesitant due to the cost of longer-term use. Credibility, legitimacy, and complementarity may well have come into play in the spreading processes. While testing the three mechanisms is beyond the scope of our paper, we provide our qualitative interpretation:

- Credibility: Typically, text advertising messages often lack credibility. Receiving the same message from multiple connected contacts can make it seem more credible.
- Legitimacy: The fact that friends had activated the data voucher might have made the then-unpopular mobile Internet look more valid and legitimate and thus helped lower the psychological barriers of adoption.
- Complementarity: During the time of our experiment, mobile messaging applications, notably Facebook, were on the rise. A big incentive to use such apps is their limitless communication potential, avoiding SMS message's count-based rates and word limits. The value of a communication technology depends heavily on the number of friends who also use it and can now be communicated with (i.e., there are strong network effects).

Our paper complements the rich study of contagion by introducing a rigorous framework for testing higher-order social influence that can flexibly account for both positive as well as negative (dampening) social reinforcement (e.g., spreading a rumor may become less satisfying if many people already know it). This framework rests on the idea of comparing observed individual-level adoption against a "no social reinforcement" baseline as a sharp null hypothesis. This criterion suggests that there are signs of complex contagions as long as the returns are higher (due to social reinforcement) than the decreasing returns of simple contagion even if the population-level diffusion curve is not super-exponential. Prior studies were either too lenient when considering a "no effect of reinforcing signals" null hypothesis (e.g., Centola, 2010), too strict when testing against "super-linear" or "non-submodular" functional forms derived from the threshold model, or may have missed positive social reinforcement by considering better spread (faster and/or further) on clustered networks (compared to random networks; Wan et al., 2024).[10]

Lastly, our embeddedness finding highlights the role of the local network structure that supports or constrains social reinforcement: the more friends two individuals have in common—the more embedded their relationship is—the stronger the effect. Theoretically, embedded relationships may represent structurally robust peers that support the confirmation, legitimacy, and credibility needed to adopt costly behaviors. More broadly in sociological field experiments, the embeddedness finding demonstrates the utility of "digital field experiments" that make use of digital infrastructure and more importantly preserve existing social ties (Parigi et al., 2017; Salganik, 2019; Van de Rijt et al., 2014), which enables testing treatment heterogeneity analysis

---

[10] It turns out that "faster spread on clustered networks" is a *possible* outcome for adoption processes with positive social reinforcement but it is neither necessary nor even particularly common across a variety of different parameter settings (Wan et al., 2025).

connected to a theoretical question of what relational characteristics of social ties are effective for complex contagion.

This study is not without its limitations. First, our primary identification of causal effects is based on observing behavior spillovers: the behavior of one adopting individual affects the behavior of another individual (i.e., that person also adopts). However, we do not observe transmission between individuals and cannot precisely speak to the mechanism underlying the peer influence (e.g., changed preferences or increased utility). Second, our study used a minimal form of clustered seeding of encouragements (one vs. two targeted individuals). Hence, we are not able to draw a conclusion about the effects of higher levels of exposures on product adoption. Third, while our experiment effectively traced adoption using unique product codes in a large social network, our data does not permit heterogeneity analysis due to privacy restrictions. We are unaware of our study participants' individual attributes, socioeconomic statuses, and geographic information. An exciting future research venue will be how behavioral thresholds change depending on customer characteristics, such as whether new product adoption is easier for individuals who fare better economically.

In conclusion, our field experiment employed both design and inference strategies to enhance the causal identification of complex contagion from social reinforcement. Beyond documenting complex product adoption contagions, we demonstrate that network embeddedness is an important structural moderator that shapes the effectiveness of social reinforcement..

Acknowledgement. We thank Guillaume St-Onge, Allison Wan, and Dean Eckles for helpful comments.

REFERENCES

Albertson, B., & Lawrence, A. (2009). After the credits roll: The long-term effects of educational television on public knowledge and attitudes. *American Politics Research*, *37*(2), 275–300.

Anderson, R. M., & May, R. M. (1992). *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press.

Aral, S. (2011). Commentary—Identifying social influence: A comment on opinion leadership and social contagion in new product diffusion. *Marketing Science*, *30*(2), 217–223.

Aral, S. (2016). Networked Experiments. In Y. Bramoullé, A. Galeotti, & B. Rogers (Eds.), *The Oxford Handbook on the Economics of Networks*. Oxford University Press.

Aral, S., Muchnik, L., & Sundararajan, A. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(51), 21544–21549.

Aral, S., & Nicolaides, C. (2017). Exercise contagion in a global social network. *Nature Communications*, *8*, 14753.

Aral, S., & Walker, D. (2011). Creating Social Contagion Through Viral Product Design: A Randomized Trial of Peer Influence in Networks. *Management Science*, *57*(9), 1623–1639.

Aral, S., & Walker, D. (2014). Tie Strength, Embeddedness, and Social Influence: A Large-Scale Networked Experiment. *Management Science*, *60*(6), 1–19.

Aronow, P. M., & Carnegie, A. (2013). Beyond LATE: Estimation of the average treatment effect with an instrumental variable. *Political Analysis*, *21*(4), 492–506.

Aronow, P. M., Eckles, D., & Peyton, K. (2018). *Icsw: Inverse Compliance Score Weighting. R package version 1.0.0*. https://cran.r-project.org/package=icsw

Athey, S., Eckles, D., & Imbens, G. W. (2018). Exact p-Values for Network Interference. *Journal of the American Statistical Association*, *113*(521), 230–240.

Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. (2012). The Role of Social Networks in Information Diffusion. *Proceedings of the 21st International Conference on World Wide Web*, *1201.4145*, 519–528.

Blume, L. E. (1993). The Statistical Mechanics of Strategic Interaction. *Games and Economic Behavior*, *5*(3), 387–424.

Bradlow, E. (1998). Encouragement designs: An approach to self-selected samples in an experimental design. *Marketing Letters*, *9*, 383–391.

Burt, R. S. (1987). Social contagion and innovation: Cohesion versus structural equivalence. *American Journal of Sociology*, *92*(6), 1287–1335.

Centola, D. (2010). The spread of behavior in an online social network experiment. *Science*, *329*(5996), 1194–1197.

Centola, D. (2018). *How behavior spreads: The science of complex contagions*. Princeton University Press.

Centola, D., & Macy, M. (2007). Complex Contagions and the Weakness of Long Ties. *American Journal of Sociology*, *113*(3), 702–734.

Christakis, N. A., & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New England Journal of Medicine*, *357*(4), 370–379.

Coleman, J. S. (1988). Social Capital in the Creation of Human Capital. *American Journal of Sociology*, *94*(s1), S95–S120.

Coleman, J. S., Katz, E., & Menzel, H. (1966). *Medical Innovation: A Diffusion Study*. The Bobbs-Merrill Company.

Conley, T. G., & Udry, C. R. (2010). Learning about a new technology: Pineapple in Ghana. *American Economic Review*, *100*(1), 35–69.

DiMaggio, P., & Garip, F. (2012). Network effects and social inequality. *Annual Review of Sociology*, *38*, 93–118.

Dodds, P. S. (2018). Slightly generalized contagion: Unifying simple models of biological and social spreading. *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks*, 67–80.

Dodds, P. S., & Watts, D. J. (2005). A generalized model of social and biological contagion. *Journal of Theoretical Biology*, *232*(4), 587–604.

Duflo, E., & Saez, E. (2003). The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Quarterly Journal of Economics*, *118*(3), 815–842.

Eagle, N., Pentland, A., & Lazer, D. (2009). Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(36), 15274–15278.

Easley, D., & Kleinberg, J. (2010). *Networks, Crowds, and Markets*. Cambridge University Press.

Eckles, D., Kizilcec, R. F., & Bakshy, E. (2016). Estimating peer effects in networks with peer encouragement designs. *Proceedings of the National Academy of Sciences*, *113*(27), 7316–7322.

Eckles, D., Mossel, E., Rahimian, M. A., & Sen, S. (2024). Long ties accelerate noisy threshold-based contagions. *Nature Human Behaviour*, *8*, 1057–1064.

Ferrali, R., Grossman, G., Platas, M. R., & Rodden, J. (2020). It takes a village: Peer effects and externalities in technology adoption. *American Journal of Political Science*, *64*(3), 536–553.

Fink, C., Schmidt, A., Barash, V., Cameron, C., & Macy, M. (2016). Complex contagions and the diffusion of popular Twitter hashtags in Nigeria. *Social Network Analysis and Mining*, *6*(1), 1–19.

Fisher, R. A. (1925). *Statistical Methods for Research Workers*. Oliver and Boyd.

Galeotti, A., Goyal, S., Jackson, M. O., Vega-Redondo, F., & Yariv, L. (2010). Network games. *Review of Economic Studies*, *77*(1), 218–244.

Gao, J., Ghasemiesfeh, G., Jones, J. J., & Schoenebeck, G. (2019). Complex Contagions in Charitable Donations. *SocArXiv*.

Gelper, S., Lans, R. van der, & Bruggen, G. van. (2021). Competition for Attention in Online Social Networks: Implications for Seeding Strategies. *Management Science*, *67*(2), 1026–1047.

Gerber, A. S., & Green, D. P. (2012). *Field experiments: Design, analysis, and interpretation*. W.W. Norton.

Godes, D., & Mayzlin, D. (2009). Firm-created word-of-mouth communication: Evidence from a field test. *Marketing Science*, *28*(4), 721–739.

Godes, D., Mayzlin, D., Chen, Y., Das, S., Dellarocas, C., Pfeiffer, B., Libai, B., Sen, S., Shi, M., & Verlegh, P. (2005). The firm's management of social interactions. *Marketing Letters*, *16*(3–4), 415–428.

Goldenberg, J., Libai, B., & Muller, E. (2001). Talk of the Network: A Complex Systems Look at the Underlying Process of Word-of-Mouth. *Marketing Letters*, *12*(3), 211–223.

Granovetter, M. S. (1973). The Strength of Weak Ties. *American Journal of Sociology*, *78*(6), 1360–1380.

Granovetter, M. S. (1978). Threshold Models of Collective Behavior. *American Journal of Sociology*, *83*(6), 1420–1443.

Granovetter, M. S. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, *91*(3), 481–510.

Guilbeault, D., & Centola, D. (2021). Topological measures for identifying and predicting the spread of complex contagions. *Nature Communications*, *12*(1), 4430.

Hébert-Dufresne, L., Scarpino, S. V., & Young, J.-G. (2020). Macroscopic patterns of interacting contagions are indistinguishable from social reinforcement. *Nature Physics*, *16*(4), 426–431.

Hinz, O., Skiera, B., Barrot, C., & Becker, J. U. (2011). Seeding strategies for viral marketing: An empirical comparison. *Journal of Marketing*, *75*(6), 55–71.

Horrace, W. C., & Oaxaca, R. L. (2006). Results on the bias and inconsistency of ordinary least squares for the linear probability model. *Economics Letters*, *90*(3), 321–327.

Jackson, M. O. (2019). The friendship paradox and systematic biases in perceptions and social norms. *Journal of Political Economy*, *127*(2), 777–818.

Johnson, C., Dowd, T. J., & Ridgeway, C. L. (2006). Legitimacy as a social process. *Annual Review of Sociology*, *32*, 53–78.

Karsai, M., Iñiguez, G., Kaski, K., & Kertész, J. (2014). Complex contagion process in spreading of online innovation. *Journal of the Royal Society Interface*, *11*(101), 20140694.

Katona, Z., Zubcsek, P. P., & Sarvary, M. (2011). Network effects and personal influences: The diffusion of an online social network. *Journal of Marketing Research*, *48*(3), 425–443.

Katz, E., & Lazarsfeld, P. F. (1955). *Personal Influence* (Free Press).

Katz, M. L., & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, *75*(3), 424–440.

Kumar, V., & Sudhir, K. (2021). Can Random Friends Seed More Buzz and Adoption? Leveraging the Friendship Paradox. *Management Science*, in press.

Lee, J. Y., & Bell, D. R. (2013). Neighborhood social capital and social learning for experience attributes of products. *Marketing Science*, *32*(6), 960–976.

Manski, C. F. (2000). Economic analysis of social interactions. *Journal of Economic Perspectives*, *14*(3), 115–136.

Mayzlin, D. (2006). Promotional chat on the internet. *Marketing Science*, *25*(2), 155–176.

Moe, W. W., & Trusov, M. (2011). The value of social dynamics in online product ratings forums. *Journal of Marketing Research*, *48*(3), 444–456.

Moffitt, R. A. (2001). Policy Interventions, Low-Level Equilibria, and Social Interactions. In S. N. Durlauf & H. P. Young (Eds.), *Social Dynamics* (pp. 45–82). MIT Press.

Mønsted, B., Sapieżyński, P., Ferrara, E., & Lehmann, S. (2017). Evidence of complex contagion of information in social media: An experiment using Twitter bots. *PLOS ONE*, *12*(9), e0184148.

Moody, J., & White, D. R. (2003). Structural Cohesion and Embeddedness: A Hierarchical Concept of Social Groups. *American Sociological Review*, *68*(1), 103–127.

Parigi, P., Santana, J. J., & Cook, K. S. (2017). Online field experiments: Studying social interactions in context. *Social Psychology Quarterly*, *80*(1), 1–19.

Rajkumar, K., Saint-Jacques, G., Bojinov, I., Brynjolfsson, E., & Aral, S. (2022). A causal test of the strength of weak ties. *Science*, *377*(6612), 1304–1310.

Rishika, R., & Ramaprasad, J. (2019). The Effects of Asymmetric Social Ties, Structural Embeddedness, and Tie Strength on Online Content Contribution Behavior. *Management Science*, *65*(7), 3398–3422.

Rogers, E. M. (2003). *Diffusion of Innovations* (5th ed.). Free Press.

Rohlfs, J. (1974). A theory of interdependent demand for a communications service. *Bell Journal of Economics and Management Science*, *5*(1), 16–37.

Salganik, M. (2019). *Bit by bit: Social research in the digital age*. Princeton University Press.

Shalizi, C. R., & Thomas, A. C. (2011). Homophily and Contagion Are Generically Confounded in Observational Social Network Studies. *Sociological Methods & Research*, *40*(2), 211–239.

St-Onge, G., Iacopini, I., Latora, V., Barrat, A., Petri, G., Allard, A., & Hébert-Dufresne, L. (2022). Influential groups for seeding and sustaining nonlinear contagion in heterogeneous hypergraphs. *Communications Physics*, *5*(1), 25.

St-Onge, G., Sun, H., Allard, A., Hébert-Dufresne, L., & Bianconi, G. (2021). Universal nonlinear infection kernel from heterogeneous exposure on higher-order networks. *Physical Review Letters*, *127*(15), 158301.

Strang, D., & Soule, S. A. (1998). Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills. *Annual Review of Sociology*, *24*(1), 265–290.

Tellis, G. J., MacInnis, D. J., Tirunillai, S., & Zhang, Y. (2019). What Drives Virality (Sharing) of Online Digital Content? The Critical Role of Information, Emotion, and Brand Prominence. *Journal of Marketing*, *83*(4), 1–20.

Ugander, J., Backstrom, L., Marlow, C., & Kleinberg, J. (2012). Structural diversity in social contagion. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(16), 5962–5966.

Uzzi, B. (1996). The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect. *American Sociological Review*, *61*(4), 674–698.

Uzzi, B. (1997). Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness. *Administrative Science Quarterly*, *42*(1), 35–67.

Valente, T. W. (1995). *Network Models of the Diffusion of Innovations*. Hampton Press.

Van de Rijt, A., Kang, S. M., Restivo, M., & Patil, A. (2014). Field experiments of success-breeds-success dynamics. *Proceedings of the National Academy of Sciences*, *111*(19), 6934–6939.

Walker, D., & Muchnik, L. (2014). Design of randomized experiments in networks. *Proceedings of the IEEE*, *102*(12), 1940–1951.

Wan, A., Riedl, C., & Lazer, D. (2025). Diffusion of Complex Contagions is Shaped by a Tradeoff Between Reach and Reinforcement. *Proceedings of the National Academy of Sciences*, *122*(28) e2422892122.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of "small-world" networks. *Nature*, *393*(6684), 440–442.

Young, H. P. (2011). The dynamics of social innovation. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(Supplement 4), 21285–21291.

Appendix for
Complex Contagion in Social Networks: Causal Evidence from a Country-Scale Field Experiment

CONTENTS

1. Conceptual Model

To illustrate the infection force equation given in Eq. (1) we show an example using $\beta = 0.1$ (i.e., 10% infection likelihood) in Figure A1.

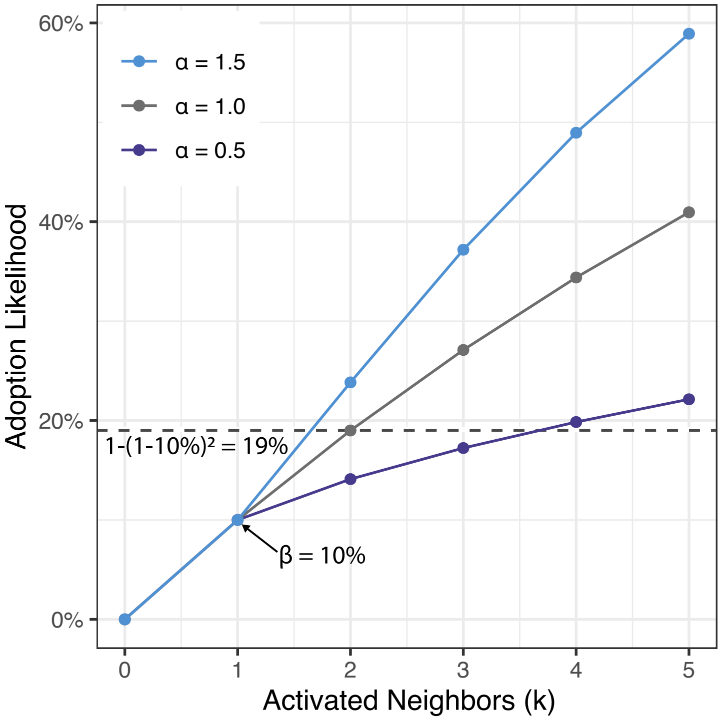

Figure A1. Example illustration of contagion at $\beta = 0.1$.

We leverage the tunable social reinforcement parameter (α) in our infection force equation given in Eq. (1) and show that in most realistic settings, "complex" contagion processes (those that feature positive social reinforcement) do not necessarily exhibit superexponential spreading (i.e., they look quite similar to "simple" contagion processes; Figure A2).

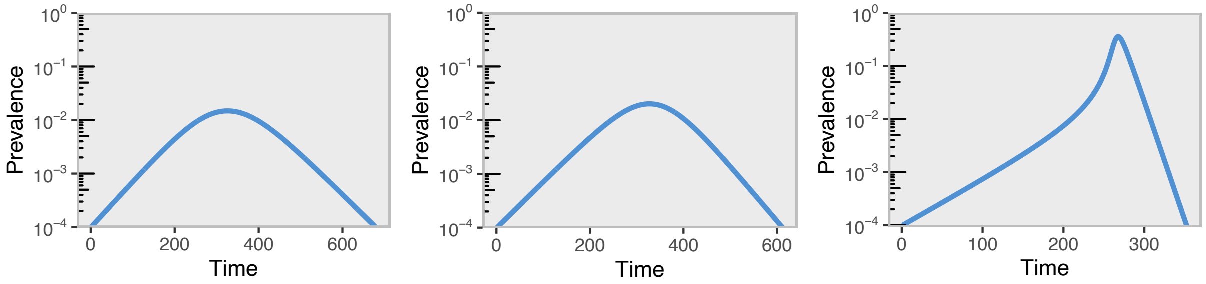

Figure A2. Across a wide range of realistic parameters (R_0 close to 1) the diffusion processes lead to standard exponential spreading (and not superexponential) despite the presence of social reinforcement ($\alpha > 1$). Left: standard exponential spreading of a simple contagion process without social reinforcement ($R_0 = 1.2$ and $\alpha = 1.0$); Middle: visually similar standard exponential spreading despite positive social reinforcement ($R_0 = 1.2$ and $\alpha = 1.2$). Right: superexponential spreading arises only when contagiousness and/or social reinforcement are very high ($R_0 = 1.2$ and $\alpha = 2$).

We simulate the spread of a behavior in a susceptible population using a standard Susceptible-Infected-Recovered (SIR) epidemic process (Barrat et al., 2008) on a random network (random regular graph with $\langle k = 20\rangle$). Under the assumptions of our model, it is almost impossible to observe superexponential growth in most realistic settings (near criticality when the basic reproduction number (R_0) describing contagiousness or transmissibility of infections is close to 1) even when social reinforcement is strong. The intuition behind this result is that exposure to multiple infected neighbors is rare when the prevalence of the infection in the population is low and there are only few instances in which social reinforcement can unfold. Conversely, when prevalence in the population is high due to high contagiousness, even moderate levels of social reinforcement make little difference on the diffusion process and are not sufficient to push it from the standard exponential regime to a superexponential regime. Taken together, our simulation analysis shows that it is possible that both aspects can hold true at the same time: positive social reinforcement acts as a behavioral mechanism on the individual level (i.e., $\alpha > 1$), while the population-level diffusion process exhibits none of the characteristics that have been theorized—and variably been used as assessment criteria—of complex contagion such as superexponential spreading. This illustrates that individual-level mechanistic insights do not directly translate into population-level empirical criteria nor do population-level observations of the standard exponential spread necessarily rule out positive social reinforcement.

2. Variable Description

Our outcome variable is whether or not a focal node adopted (redeemed) a data voucher. The key predictor variables are Treated Neighbors (K_{ITT}) and Activated Neighbors (K), the number of a focal node's immediate neighbors that were seeded and adopted the voucher code. A core moderator variable in our study is structural embeddedness—the number of friends that two individuals have in common normalized by the degree of the two individuals. Counting common neighbors in a large network like ours with 48M nodes poses a significant computational challenge and is only feasible with the use of advanced graph algorithms. We built on the work by (An et al., 2019) to efficiently compute edge-embeddedness in our large network. We also consider individual and dyadic network characteristics that might account for heterogeneity in peer effects and the propensity of adoption. Degree is the number of the connected neighbors via calls and SMS during the three-month period. Provided that peer influence is more likely and stronger when individuals are close to each other, we control for dyadic tie weights in terms of duration and frequency of their interactions. We turned them into normalized individual-level variables in ways that their value indicates tie strength to the treated seed, relative to other network neighbors connected to that seed.

Variable	**Description**
Adoption (Y)	Outcome variable, 1 if individual adopts mobile internet by redeeming a voucher code, 0 otherwise. Unless noted otherwise, we consider the entire experiment period of 7 days.
Treated Neighbors (K_{ITT})	Number of treated neighbors that a focal node has.
Activated Peers (K)	Number of immediate neighbors that were seeded and adopted the voucher code. A focal node's seed peer is only considered "activated" if the seed adopted before the focal node adopted.
Degree	Individuals' number of friends in the baseline social network, based on calls and SMS, in a three-month period
Duration	Tie strength between the focal node and the encouraged seed node in the baseline social network measured in voice minutes that an individual had with a treated seed, relative to other network neighbors (i.e., row-normalized to the unit interval [0, 1]). For focal nodes connected to two seeds, we average the tie strength.
Frequency	Tie strength between the focal node and the encouraged seed node in the baseline social network measured as contact frequency, relative to other network neighbors (i.e., row-normalized to the unit interval [0, 1]). We use a composite measurement of (voice+SMS) - (voice×SMS), which showed a better fit than when they are included separately. For focal nodes connected to two seeds, we average the tie strength.
Embeddedness	The number of common friends that an individual had with a treated seed (M_{ij}) normalized by the degree of the focal node (D_i) and the seed (D_j): $\frac{M_{ij}}{D_i+D_j-M_{ij}-2}$. Averaged for focal nodes who are connected to both seeds.

Table A1. Variable Description.

	Mean	SD	Min	Max	(1)	(2)	(3)	(4)	(5)	(6)
Adopted Product (1)	0.00	0.05	0.00	1.00						
k_{ITT} (2)	0.71	0.47	0.00	2.00	0.01***					
k (3)	0.08	0.28	0.00	2.00	0.05***	0.17***				
Degree (4)	64.66	68.46	1.00	4513.00	0.00***	0.00***	-0.03***			
Duration (row normalized) (5)	0.05	0.13	0.00	1.00	-0.01***	-0.01***	-0.01***	-0.22***		
Frequency (row normalized) (6)	0.07	0.17	0.00	1.00	0.00	0.00***	0.02***	-0.23***	0.66***	
Embeddedness (7)	0.02	0.02	0.00	0.47	0.03***	0.03***	0.05***	-0.04***	-0.02***	0.00***

$^{***}p < 0.001$

Table A2. Descriptive statistics of study variables.

3. Contamination Issue

For the data analysis, we removed “contaminated” nodes for which we could not cleanly attribute observed effects to assigned treatment conditions. That is, we removed seed nodes who were themselves connected to other seed nodes, focal nodes who were connected to contaminated seed nodes, and focal nodes who were connected to more than two seed nodes (about 4% of nodes were affected from such contamination). One concern may be that this results in imbalanced observables (e.g., with regard to node degree). We believe this does not constitute a major problem at least with regard to estimating LATE as the majority of nodes are removed because they are connected to contaminated seeds. Since seeds are randomly assigned to treatment, this removal should not result in imbalanced observables.

The biggest challenge in implementing this marketing campaign was online-based diffusion (see [citation removed for blind peer review]). That is, a few customers posted their voucher codes to an online community and those were accessed individually across the population. While this online sharing further helped the product go viral, it certainly raised a concern for our inference of peer-to-peer diffusion through the mobile social network. However, because all seeded vouchers contained unique codes, we were able to identify the four “contaminated” codes reported with an atypically high number of adoptions (10~50 times higher activation counts than other codes). Note that this study does not aim to investigate diffusion processes in the whole population; we are interested in a specific subpopulation that comprises our experimental treatment groups—only the seeds and their one-hop neighbors who would be early spreaders and therefore unaffected by contamination. To provide conservative estimates and avoid counting false-positive product adoptions (i.e., a seed and his neighbor independently adopted the voucher from the online source, but they happened to have a tie to each other), we excluded the cases where any of the two nodes in the randomly assigned edge had adopted a contaminated voucher code. This measure reduced our sample size, but not dramatically (2,640,711 → 2,502,187, a 5% decrease). In a similar vein, we only counted non-contaminated adoption codes.

4. Intent-to-Treat Analysis

Here we consider treatment effect heterogeneity (Table A3).

Dependent Variable:	**Adopted Product within first ...**	
	6h	7 days
	(1)	(2)
Intercept	0.0001^{***}	0.0013^{***}
	(0.0000)	(0.0001)
$k_{ITT} = 1$	$-0.0000^{\dagger}$	0.0000
	(0.0000)	(0.0001)
$k_{ITT} = 2$	0.0005	0.0023^{*}
	(0.0003)	(0.0009)
Embeddedness		
× $k_{ITT} = 1$	0.0134^{***}	0.0371^{***}
	(0.0015)	(0.0036)
× $k_{ITT} = 2$	$0.0132^{\dagger}$	0.0495^{**}
	(0.0076)	(0.0190)
Controls		
Embeddedness	0.0027^{***}	0.0248^{***}
	(0.0008)	(0.0026)
Degree	-0.0000^{**}	-0.0001^{***}
	(0.0000)	(0.0000)
Duration	-0.0005^{***}	-0.0035^{***}
	(0.0001)	(0.0002)
Frequency	0.0005^{***}	0.0015^{***}
	(0.0001)	(0.0002)
Adj. R^2	0.0005	0.0010
Num. obs.	2,502,187	

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{\dagger}p < 0.1$

Table A3. Regression results of ITT treatment effect heterogeneity.

Notice that spreading along indirect paths cannot explain these results. To see this, it is useful to recall that two-hop spreading along an indirect path happens only at the squared likelihood of the direct path as both nodes along the indirect paths would have to activate (i.e., $1 - (1 - \beta^2)^N$ where N is the number of available indirect paths). Empirically, we observe seed nodes have about 2.8 friends in common with focal nodes so that this adoption likelihood is about two orders of magnitude *smaller* than the main effect (see the conceptual Figure A1 above). However, empirically, we find a coefficient that is instead about one order of magnitude *larger* than the main effect. This supports the conclusion that our results are inconsistent with the hypothesis that spreading is driven by indirect paths and instead much more consistent with the hypothesis that the tie strength of structurally embedded relationships explains the effect. Finally, notice also that in the model using adoption within the first 6h indirect spreading is effectively ruled out as shown by the near-zero the intercept which indicates that there simply is virtually no indirect spreading at all.

We illustrate the potential effect that indirect spreading along two-hop paths may have relative to the effect of direct spreading in Figure A3. Indirect spreading has only a very small effect relative to the direct effect, especially given the small number of 2.8 indirect paths that we observe in our data between focal nodes and seed nodes (so a total of 5.6 indirect paths between the focal node and the two seeds).

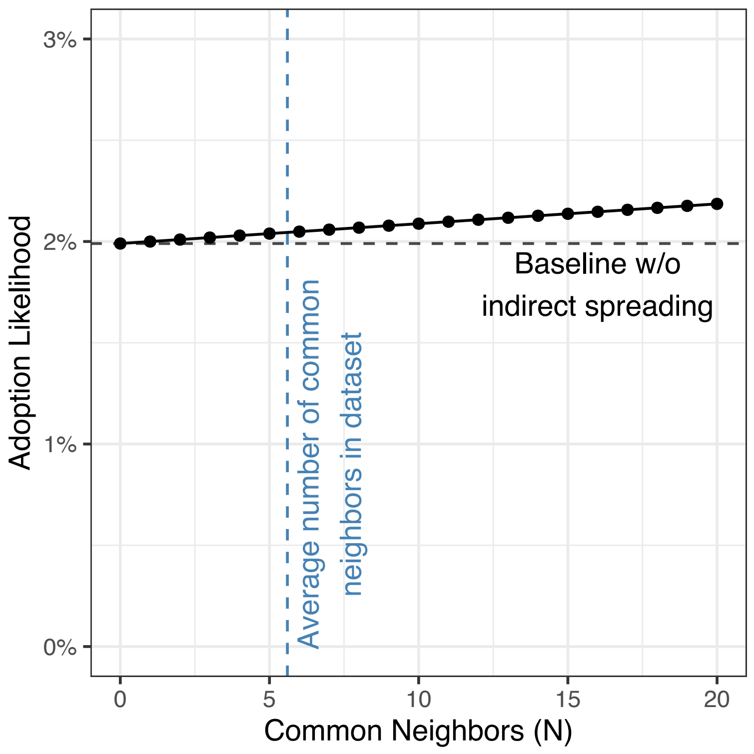

Figure A3. Example illustration of effect of indirect spreading at $\beta = 0.01; k = 2$ and $\alpha = 1$ (simple contagion).

5. Analysis of Encouragement Design with 2SLS

Conditions Here we consider several conditions in applying the 2SLS IV analysis for the peer encouragement design. The potential outcome formulation of the encouragement design requires two assumptions (Eckles et al., 2016). First, it requires an exclusion restriction for instrumental variables (IVs). The exclusion restriction seems plausible in our case as participants are blinded to their assignment. It seems safe to assume that participants did not realize they were randomly assigned to treatments—especially given the minimal nature of the encouragement—or if they were it was just a normal kind of marketing campaign that would not otherwise affect their behavior. The minimal nature of the encouragement makes it plausible that it affects individuals only by causing them to have contact with an activated neighbor and their neighbors own use of mobile internet; however, there may be effects of the encouragement not captured by the adoption status of seed nodes (e.g., changes in content of communication). Next, is the no interference assumption and direct-effect-bounded interference assumption. We expect this assumption to be violated in this setting, even in our finite population. To limit the threat of interference we focus on only one-hop neighbors, exclude individuals who are connected to multiple seed pairs, and the fact that observations in our experiment make up only a relatively small portion of the entire network (about 3M nodes in 48M network). Overall, we therefore work under the assumption that the remaining interference is small compared with the effects of interest (see Eckles et al. 2016 for more details and why this assumption may be justified). To address remaining concerns, we present a complementary analysis using a Fisher exact p-value test that is specifically designed to address the issue of network interference (see the method description in 3.2.3 and its findings in 4.3.). The 2SLS estimate using binary treatment indicators formed from the levels of the encouragement Z_i, estimates a weighted average of effects of changes in increments of D_i using a single binary indicator (Imbens & Angrist, 1994), each of which estimates an average causal response (ACR).

With the parameter estimates for discrete levels of K_i from our 2SLS analysis of the encouragement design in hand, we can then proceed to test our hypothesis of positive reinforcement from redundant exposure. Intuitively, we take the parameter estimates of having zero, one, or two activated neighbors and then apply our criterion for complex contagion to test if the increase in adoption likelihood from exposure to two activated neighbors is larger than the expected counterfactual from two activated neighbors under simple contagion. Finally, we then use the same estimation framework to test for a moderating effect of embeddedness. Specifically, we will add embeddedness as a covariate to the first-stage regression and an interaction term between K_i and embeddedness to the second-stage regression.

First-stage Results

Dependent Variable:	**One Activated Seed Neighbor**		**Two Activated Seed Neighbor**	
	W/O Controls	W/ Controls	W/O Controls	W/ Controls
Intercept	0.064***	0.056***	−0.000	−0.000**
	(0.000)	(0.000)	(0.000)	(0.000)
Single exposure: tie with both seeds	0.050***	0.039***	0.017***	0.017***
	(0.002)	(0.002)	(0.000)	(0.000)
Multiple exposure: tie with one seed	0.047***	0.047***	0.000	−0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Multiple exposure: tie with both seeds	0.119***	0.108***	0.050***	0.049***
	(0.002)	(0.002)	(0.000)	(0.000)
Degree		−0.009***		−0.000***
		(0.000)		(0.000)
Duration		−0.082***		−0.001***
		(0.002)		(0.000)
Frequency		0.063***		0.000***
		(0.001)		(0.000)
Embeddedness		0.371***		0.004***
		(0.007)		(0.001)
R^2	0.008	0.011	0.037	0.037
Adj. R^2	0.008	0.011	0.037	0.037
Num. obs.	2502187	2502187	2502187	2502187
F statistic	6488.094	3884.501	31677.524	13593.571

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A4. First-stage model results. Omitted category is "Single exposure: tie with one seed."

Dependent Variable:	Adopted Product within first ...					
	6h	12h	24h	36h	48h	7 days
	(1)	(2)	(3)	(4)	(5)	(6)
Peer Effects						
Intercept	-0.000^{**}	-0.000	0.000	0.000	0.000	0.001^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$k = 1$	0.002^{***}	0.003^{***}	0.004^{***}	0.005^{***}	0.005^{***}	0.010^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$k = 2$	0.011^{*}	0.015^{*}	0.016^{*}	0.023^{**}	0.026^{**}	0.060^{***}
	(0.005)	(0.006)	(0.007)	(0.008)	(0.008)	(0.012)
Controls						
Degree	0.000	0.000	0.000	0.000	0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Duration	-0.000^{***}	-0.001^{***}	-0.001^{***}	-0.001^{***}	-0.001^{***}	-0.003^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Frequency	0.000^{***}	0.000^{**}	0.000^{**}	0.001^{***}	0.001^{***}	0.001^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Embeddedness	0.012^{***}	0.016^{***}	0.019^{***}	0.026^{***}	0.028^{***}	0.047^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
α	2.622	2.556	2.297	2.551	2.449	3.084
Wald $P^{Obs}_{k=2} = P^{Simple}_{k=2}$	8.290	8.694	6.244	10.151	10.455	36.185
Adj. R^2	0.001	0.001	0.001	0.002	0.002	0.003
Num. obs.			2,502,187			

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A5. ATE effects of activated neighbors on product adoption, replicated in different time windows (Model 4 in Table 2).

Comparison of OLS and 2SLS Estimates. When analyzing peer effects through an instrumental variable framework, it is not uncommon to obtain different effect sizes between OLS and 2SLS estimates (Eckles et al., 2016; Aral and Nicolaides 2017). Eckles et al. (2016) suggests that it cannot be assumed that OLS will necessarily be an upper bound the true effect. This could be due to confounding, simultaneity, or the fact that these different models identify different causal quantities. In our analysis, the effect size was larger in the 2SLS model. This suggests that OLS underestimated the peer effects of *k*, accompanied by an unidentified process where the adoption likelihood *Y decreased* as the endogenous variable *k* increased.

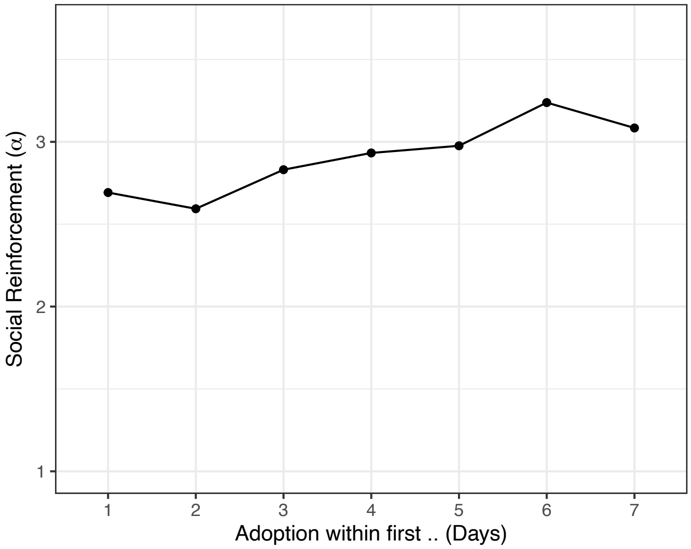

Figure A4. Comparison of the strength of social reinforcement across different time windows (ATE effects from Model 4 in Table 2 estimated on different time windows).

6. Exact P-Values for Network Interference

Network interference is a well-documented problem for networked experiments (Aral, 2016; Aral & Walker, 2014; Aronow, 2012). Randomized controlled experiments on networked populations often violate the stable unit treatment value assumption (SUTVA) because treated individuals are usually connected to other non-treated individuals (Gerber & Green, 2012; Rubin, 1980). SUTVA requires that the outcomes for any given unit be unaffected by other units' treatment assignment. This requirement serves as a condition for the unbiased inference of average treatment effects. Since our experiment was conducted on a single network, even treatments that are randomly assigned do not completely rule out a biased inference. This poses a threat to the inference of treatment effects despite the randomized design, potentially leading to inflated Type 1 errors.

In the Methods section, we have provided an overview of the exact p-value approach to estimating spillovers in Athey, Eckles, Imbens (2018, AEI henceforth). Here, we offer details of hypothesis testing that we have reported in the Results section 4.3. The null model and the form of the peer effect term differ by the specific hypotheses to test. We explain these, given N individuals i (nodes), a vector of random treatment assignments W_i, and a network adjacency matrix G with $G_{ij} = 1$ if individual i and j are connected and 0 otherwise.

Hypothesis 0 (No Spillover): We consider the null hypothesis that for the outcome Y_i, only the treatment assignment W_i is relevant, and the treatment assignment for other individuals $W_{j \neq i}$, and consequently exposure to treated neighbors, does not matter. The null hypothesis of no spillover effects is that outcome of Y_i is only affected by the treatment of w_i but not neighbor's treatment (this is H2 in the original AEI). Formally, this can be expressed as: $Y_i(w) = Y_i(w')$ for all i, and for all pairs of assignment vectors $w, w' \in W$ such that $w_i = w_i'$. The null model includes only the own treatment:

$$Y_i^{obs} = \alpha_0 + \tau_{treat} \cdot W_i + \epsilon_i \qquad \text{(A1)}$$

Test statistic. To assess the degree to which the observed effects are distinct from the artificial experiment, we calculate a score test statistic, T_{score}. T_{score} for H0 is the covariance between the residual of the null model and i's number of treated neighbors (this is adapted from Eq 7 in AEI). This, and all other T_{score} are calculated for all nodes that are not isolates, which is true for all nodes by design in our setting:

$$T_{score} = cov(\epsilon_i, \textstyle\sum_{j=1}^{N} W_j \cdot G_{ij}) \qquad \text{(A2)}$$

Substantially, T_{score} allows us to identify the degree to which the peer effects variable explains the leftover unexplained variance in the treatment-only model. A high T_{score} means that the error term of the no-spillover model substantially covaries with the (observed) peer effect—the number of treated neighbors in this case—and hence substantially explains adoption behavior.

Hypothesis 1 (Threshold Peer Effects): We consider the null hypothesis that for the outcome of Y_i it only matters whether i has any treated neighbors. Here we test the alternative hypothesis that the number of treated neighbors matters. The null hypothesis of threshold peer effects is that the effect of having two or more treated neighbors is not different from the effect of having one treated neighbor on i's adoption (H8 in AEI). Formally: $Y_i(w) = Y_i(w')$ for all i, and for all pairs of assignment vectors $w, w' \in W$ such that $1\{\sum_{j=1}^{N} w_j * G_{ij} > 0\} = 1\{\sum_{j=1}^{N} w_j' * G_{ij} > 0\}$. The null model includes the own treatment term and the discrete variable of $k \geq 1$:

$$Y_i^{obs} = \alpha_0 + \tau_{treat} \cdot W_i + \tau_{peer} \cdot \mathbf{1}\left\{\sum_{j=1}^{N} W_j \cdot G_{ij} \geq 1\right\} + \epsilon_i \quad \text{(A3)}$$

Test statistic. T_{score} is the covariance between the residual of the null model and i's number of treated neighbors:

$$T_{score} = cov(\epsilon_i, \sum_{j=1}^{N} W_j \cdot G_{ij}) \qquad \text{(A4)}$$

Hypothesis 2 (Peer Effect Heterogeneity): Under the null hypothesis of no peer effect heterogeneity it does not matter which specific peers of individual i received the treatment. Here we test specifically the alternative hypothesis that embeddedness matters for the effect of individual i on j. A null hypothesis of peer effect heterogeneity is that only the number of treated peers matter, not which neighbor is treated. In our case, we are specifically interested in the role of embeddedness, so we test the null hypotheses that the embeddedness—the number of common friends—that i has with a treated peer does not matter (H7 in AIE). Formally: $Y_i(w) = Y_i(w')$ for all i, and for all pairs of assignment vectors $w, w' \in W$ such that $\sum_{j=1}^{N} w_j * G_{ij} = \sum_{j=1}^{N} w_j' * G_{ij}$. The null model includes the own treatment term and the continuous peer effects variable (the number of treated neighbors of i):

$$Y_i^{obs} = \alpha_0 + \tau_{treat} \cdot W_i + \tau_{peer} \cdot \sum_{j=1}^{N} W_j \cdot G_{ij} + \epsilon_i \qquad \text{(A5)}$$

Test statistic. T_{score} is the covariance between the residual of the null model and i's number of common friends with the treated neighbor, derived from a weighted matrix G^{emb} where $G_{i,j}^{emb}$ equals to the number of common friends between i and j (adapted from Eq. 12 in AEI):

$$T_{score} = cov(\epsilon_i, \sum_{j=1}^{N} G_{ij} \cdot (G_{i,j}^{emb} \cdot W_j)) \qquad \text{(A6)}$$

For all three hypotheses, the exact p-values (2-tailed test) can be expressed as the chance that T_{score}^{Sim} exceeds the observed T_{score}^{Obs}:

$$\widehat{p - value} = 2 * \frac{1}{R} \sum_{r=1}^{R} 1\{|T_{score,r}^{Sim}| > |T_{score}^{Obs}|\}$$

where R is the total number of simulation trials (Taylor & Eckles, 2018).

Permutation Procedures: We first constructed $\mathbb{P}_F$—a subset of nodes that are fixed with their own treatment status and used to produce score test statistics—by selecting the half of nodes in the following steps: (i) randomly choose one of the seeds in the seed pair; and (ii) choose all focal nodes connected to the randomly chosen seed. The p-values are valid irrespective of the choice of $\mathbb{P}_F$, but maximizing edge comparisons by selecting one seed per the seed pair increases the power of the testing procedure (AEI: p. 235).

In each simulation trial, we permuted treatment status of those not in the fixed subset $\mathbb{P}_F$ and calculated the peer-effects term (e.g., the number of treated neighbors in **H0** and **H1**). We shuffled treatment status *only among seed nodes* to prevent the cases where focal nodes—they are by design "untreated" neighbors of seeds—become treated in the simulation. Substantively, in our simulation, a treated node in the single exposure condition may become an un-treated node, and a pair of seed nodes in the multiple exposure condition may become a single exposure pair. Focal nodes cannot become seed nodes. The distribution of simulated test statistics is shown in Figure A5.

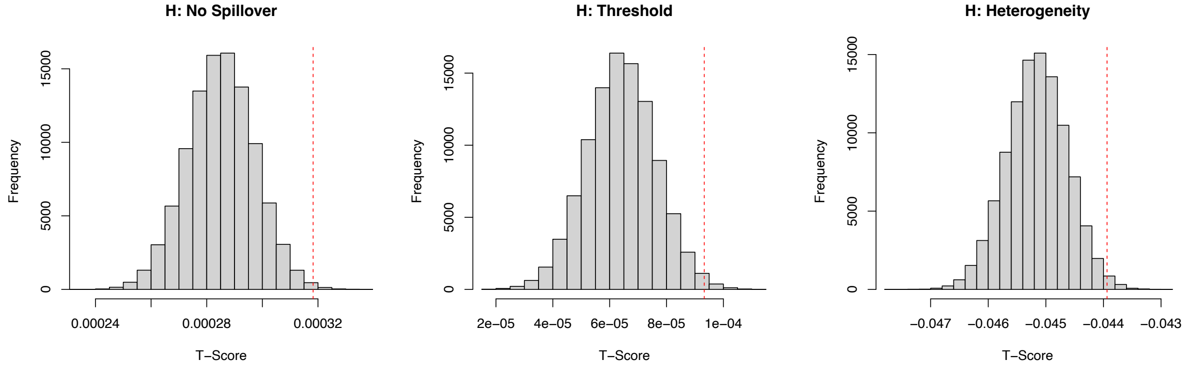

Figure A5. The distribution of test statistics from the exact *p*-value method testing H0: Spillover, H1: Threshold Peer Effects, and H2: Peer Effect Heterogeneity (by embeddedness).

7. Related Literature

Methodologically, our experiment stands as the first large-scale field experiment outside online social media that randomizes the essential theoretical element in the discussion of complex contagions. We put the features of our experiment in perspective by surveying what other studies have contributed. Table A6 is a summary of how our study differs from closely related prior work. Contrary to prior studies that focused on social media (Bakshy et al., 2012), our experiment involves a single, novel product with trackable unique codes so that we can identify precisely which neighbor was more influential. This design prevents social influence from being confounded with the nature of the product which might be the case when studying the spread of such as memes on Twitter or apps on Facebook. Our study context is realistic to the extent that awareness of peer adoption was entirely up to consumers and their interactions—with no automatic notifications helped by an online system.

In terms of the causal effects of the number of exposure to activated peers, which is the central focus of our study, only two studies (Bakshy et al., 2012; Centola, 2010) provide estimates from a randomized experiment. But Bakshy et al.'s estimate was based on the observational association, as their experiment only randomized exposure or non-exposure. Centola's experiment was based on a randomized network structure but in an artificial network that eliminated existing social ties. In Figure A6, we plotted the relative rates of adoptions when there are k and k-1 adopting neighbors, reported in the past studies and our study.

	Studies Complex Contagion?[a]	Randomized Experiment?[b]	Randomized Number of Exposures?[c]	Testing Structural Heterogeneity?[d]	Study Outcome[e]	Network Type[f]
This Study	*Yes*	*Yes*	*Yes*	*Yes*	*Product Adoption*	*Offline social network*
Centola (2010)	Yes	Yes	No	No	Product Adoption	Artificial network
Bakshy et al. (2012)	Yes	Yes	No	No	Content Sharing	Online social network
Ferrali et al. (2020)	Yes	No	No	No	Product Adoption	Offline social network
Karsai et al. (2014)	Yes	No	No	No	Product Adoption	Contact network
Ugander et al. (2012)	Yes	No	No	Yes	Product Adoption	Online social network
Aral et al. (2009)	Yes	No	No	No	Product Adoption	Online social network
Iyengar & Berger (2014)	Yes	No	No	No	Product Adoption	Online social network
Aral & Nicolaides (2017)	Yes	No	No	Yes	Behavioral Change	Online social network
State & Adamic (2015)	Yes	No	No	No	Content Sharing	Online social network
Mønsted et al. (2017)	Yes	No	No	No	Content Sharing	Online social network
Fink et al., (2016)	Yes	No	No	No	Content Sharing	Online social network
Weng et al. (2013)	Yes	No	No	Yes	Content Sharing	Online social network

Table A6. Literature summary.

Note: [a] We only included empirical studies that explicitly addressed the theory of complex contagions and provided at least a partial test of the theory.

[b] We classify as a randomized experiment when researchers imposed a treatment on randomly assigned individuals/groups to estimate peer effects.

[c] This column focuses on whether the effects of an additional exposure were found as an observational association or an exogenous experimental manipulation.

[d] This column indicates whether or not they additionally investigated network-structural variation of peer effects. For example, this study tests the moderating role of network embeddedness in inducing the effect of additional exposure. Ugander et al. (2012) found strong effects of structural diversity (number of components), less the number of exposures. Aral & Nicolaides (2017) similarly tested network embeddedness as a peer influence moderator but do not link it to the heterogeneity of the effectiveness of complex contagions.

[e] Content sharing includes the diffusion of memes and hashtags (Fink et al. 2016; Mønsted et al. 2017; Weng et al. 2013) and URL sharing (Bakshy et al. 2012).

[f] It refers to the context in which network nodes and ties were defined and interactions among peers occurred. Our study uses a whole network of mobile phone users unlike other studies based on Twitter, Facebook, or other created communities.

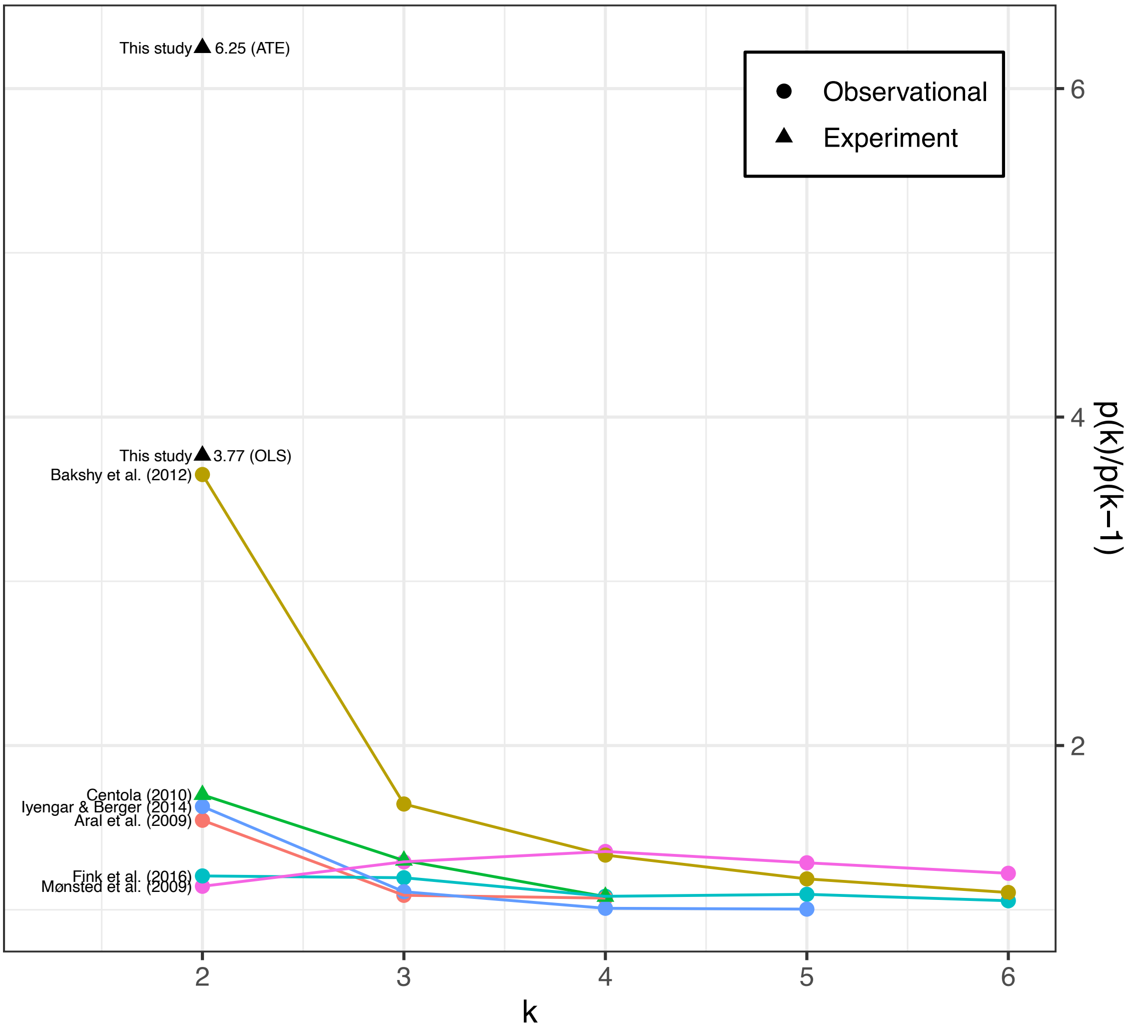

Figure A6. Adoption ratios by *k* reported in past empirical studies.
Note: This plot adapts and extends Figure S1 of Eckles et al. (2019). We selected the empirical studies that show peer effects by specific number of adopted neighbors. Note that other than Bakshy et al. (2012), there is no study showing a multiplied likelihood comparing *k*=2 and *k*=1, even when evidence of complex contagions was claimed. Bakshy et al.'s estimate was based on the observational association, as their experiment only randomized exposure or non-exposure. Values for our study are based on estimates from Table 2 Model 4 (ATE) and Model 5 (OLS) (i.e., based on the same values as shown in Figure 3). The characteristics of these studies are summarized in Table A5.

8. The Benefits of Clustered Targeting

Can a manager leverage the benefits of social reinforcement to design a marketing campaign? We examine the implications of our findings for the design of a seeding strategy in viral marketing, a topic of significant interest. The idea is that, if product adoption is substantially improved by social reinforcement engineered by multiple contacts (i.e., clustered targeting: seeding "pair" of individuals), we can leverage it to maximize early spillovers from the seeding stage.

Intuitively the effectiveness of clustered targeting will depend on (a) how likely it is that two clustered seed individuals "activate" together to expose their neighbors to social reinforcement in the first place and (b) the strength of the social reinforcement effect.[1] We construct a stylized model to illustrate how the potential benefits of clustered targeting vary with the adoption probability of seeds and the amplification of the social reinforcement effect. In the model, given adoption probability β, seeds activate to create exposure for one-hop neighbors, and these neighbors may subsequently adopt. We vary the effect of social reinforcement by multiplying the $k = 2$ effect by an amplification factor α. In turn, one-hop neighbors' resulting adoption probability is calculated by:

$$[p(exposure\ to\ one\ seed) \times \beta] + [p(exposure\ to\ both\ seeds) \times \beta \times \alpha]$$

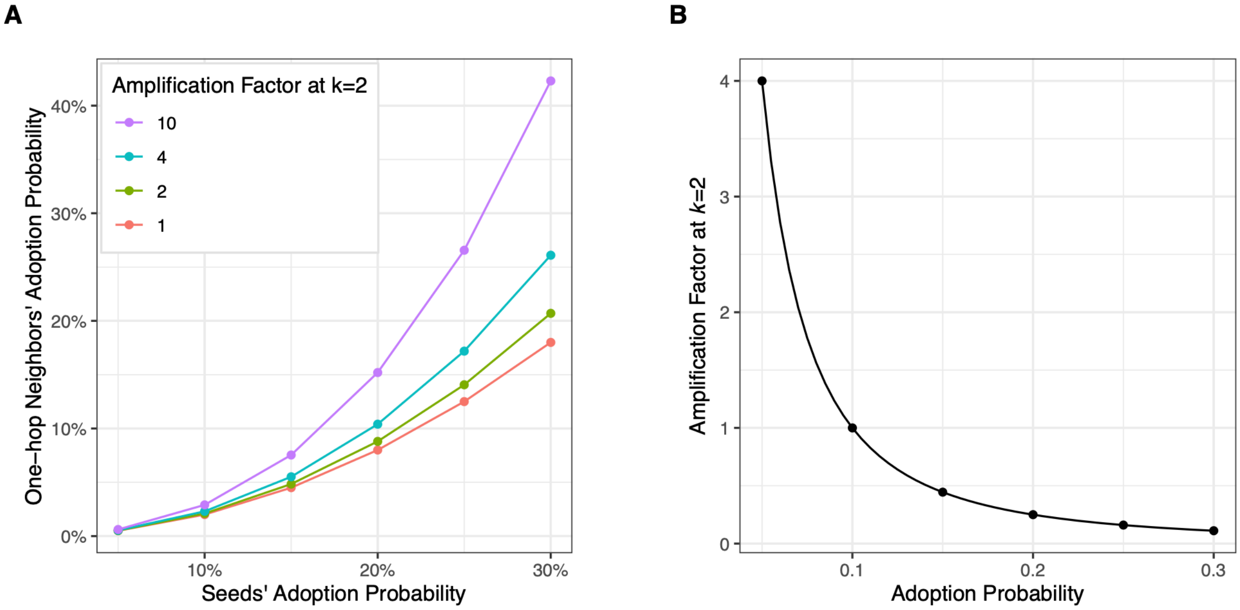

Figure A7. The potential benefits of clustered targeting as a function of seed's adoption probability and amplification of social reinforcement

The results are summarized in Figure A7-A. The key finding is that there are additional benefits generated by amplification at $k = 2$ but to a limited degree. The primary reason is that exposure to two activated peers is necessarily rare. For example, say a customer is tied to two seeds and these seeds adopt with 10% probability. Then the probability for this customer to have two activated peers is only 1% ($0.1^2 = 0.01$). This means that the proportion of adoption generated by $k = 2$ exposure would comprise a very small portion of the total adoption created by clustered targeting. We thus see little difference made by social reinforcement at $\beta = 0.1$.

The simulation result shows that the effect of social reinforcement materializes at a higher rate when a high adoption probability creates more $k = 2$ exposure cases. Yet we note other realisms that were not parameterized in this stylized model but could be detrimental to the benefits of clustered seeding. One such realism is the degree to which an individual happens to have ties to both seeds—in other words, the probability that one forms a triangle to two seeds. Among our focal nodes from the single and multiple exposure treatments, only 2.2% had ties to both seeds. This would further lower the adoption generated by $k = 2$ exposure

---

[1] Naturally, it will also depend on how many friends the two activating individuals have in common but since we are trying to optimize the marketing campaign given the network this is merely a moderating factor.

against the total adoption among clustered seeds' neighbors. Then, clustered targeting with a rare $k = 2$ exposure is unlikely to perform better than random seeding primarily relying on a $k = 1$ exposure.

Next, when a marketer has a given cost of targeting individuals and hence a target adoption rate at which the campaign is profitable, how much weight should be placed on the adoption probability of a seeded customer and social reinforcement at k=2? Let us denote Θ as the target adoption rate, α as the amplification factor at $k = 2$, and β as seed's adoption rate, then $\Theta = \alpha \times \beta^2$. We can infer that regardless of the size of the target adoption rate, the desired social reinforcement effect α is inversely related to the squared term of the desired seeds' adoption probability β. Figure A7-B illustrates this point. Note that this simplified numerical example considers only the effect of clustered seeding on one-hop neighbors. In order to more fully assess whether clustered seeding is better able to start cascades beyond one-hop neighbors it becomes critically important to consider the role of network topology (Eckles et al., 2019; Pastor-Satorras & Vespignani, 2001).

9. Consistency in Linear Probability Models

We consider a concern over using linear probability models (LPM). Horrace and Oaxacana (2006) point out that "[c]onsistency seems to be an exceedingly rare occurrence as one would have to accept extraordinary restrictions on the joint distribution of the regressors. Therefore, OLS is frequently a biased estimator and almost always an inconsistent estimator of the LPM." That is, given outcomes $Y \in \{0,1\}$ and fixed covariate vectors X, the OLS estimator is consistent and unbiased under the linear probability model if $X_i^T \beta \in [0,1]$ for all i, otherwise the OLS estimator is biased and inconsistent.

As a solution to this problem, they suggest a 'trimming' estimator approach: "[t]rimming observations whose predictions lie outside the unit interval and re-estimating the OLS model (based on the trimmed sample) may reduce finite sample bias." The procedures of this approach are as follows: (i) estimate the LPM; (ii) drop observations whose $\hat{y}$ lies outside $[0, 1]$ until $X_i^T \beta \in [0,1]$ for all i; and (iii) re-estimate the model with the remaining observations.

We adopt the trimming approach and examine whether the OLS estimates in our ITT models are consistent even after we trim observations (see Table A7). There are varying rates of observations being trimmed. The trimmed cases are more prevalent in earlier time windows (42% at 6 hours to 6% at 7 days). But, even after trimming out such cases, our findings regarding the difference of adoption probabilities between $k_{ITT} = 1$ and $k_{ITT} = 2$ generally hold, indicated by Wald test statistics large enough to reject the null hypothesis of $p_{k=2}^{Obs} = p_{k=2}^{Simple}$ and $\alpha > 1$ indicative of positive contagion effects across all the time windows.

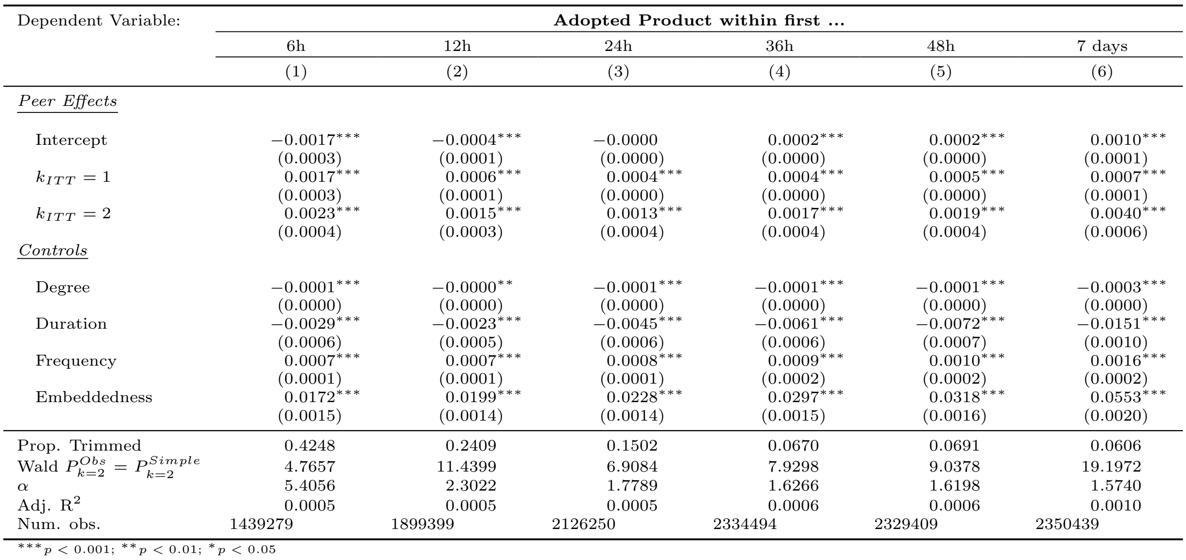

Dependent Variable:	Adopted Product within first ...					
	6h	12h	24h	36h	48h	7 days
	(1)	(2)	(3)	(4)	(5)	(6)
Peer Effects						
Intercept	-0.0017^{***} (0.0003)	-0.0004^{***} (0.0001)	-0.0000 (0.0000)	0.0002^{***} (0.0000)	0.0002^{***} (0.0000)	0.0010^{***} (0.0001)
$k_{ITT} = 1$	0.0017^{***} (0.0003)	0.0006^{***} (0.0001)	0.0004^{***} (0.0000)	0.0004^{***} (0.0000)	0.0005^{***} (0.0000)	0.0007^{***} (0.0001)
$k_{ITT} = 2$	0.0023^{***} (0.0004)	0.0015^{***} (0.0003)	0.0013^{***} (0.0004)	0.0017^{***} (0.0004)	0.0019^{***} (0.0004)	0.0040^{***} (0.0006)
Controls						
Degree	-0.0001^{***} (0.0000)	-0.0000^{**} (0.0000)	-0.0001^{***} (0.0000)	-0.0001^{***} (0.0000)	-0.0001^{***} (0.0000)	-0.0003^{***} (0.0000)
Duration	-0.0029^{***} (0.0006)	-0.0023^{***} (0.0005)	-0.0045^{***} (0.0006)	-0.0061^{***} (0.0006)	-0.0072^{***} (0.0007)	-0.0151^{***} (0.0010)
Frequency	0.0007^{***} (0.0001)	0.0007^{***} (0.0001)	0.0008^{***} (0.0001)	0.0009^{***} (0.0002)	0.0010^{***} (0.0002)	0.0016^{***} (0.0002)
Embeddedness	0.0172^{***} (0.0015)	0.0199^{***} (0.0014)	0.0228^{***} (0.0014)	0.0297^{***} (0.0015)	0.0318^{***} (0.0016)	0.0553^{***} (0.0020)
Prop. Trimmed	0.4248	0.2409	0.1502	0.0670	0.0691	0.0606
Wald $P^{Obs}_{k=2} = P^{Simple}_{k=2}$	4.7657	11.4399	6.9084	7.9298	9.0378	19.1972
α	5.4056	2.3022	1.7789	1.6266	1.6198	1.5740
Adj. R^2	0.0005	0.0005	0.0005	0.0006	0.0006	0.0010
Num. obs.	1439279	1899399	2126250	2334494	2329409	2350439

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table A7. Robustness check after repeatedly trimming observations whose $\hat{y} \notin [0,1]$

10. Algorithm to find a solution for infection force equation

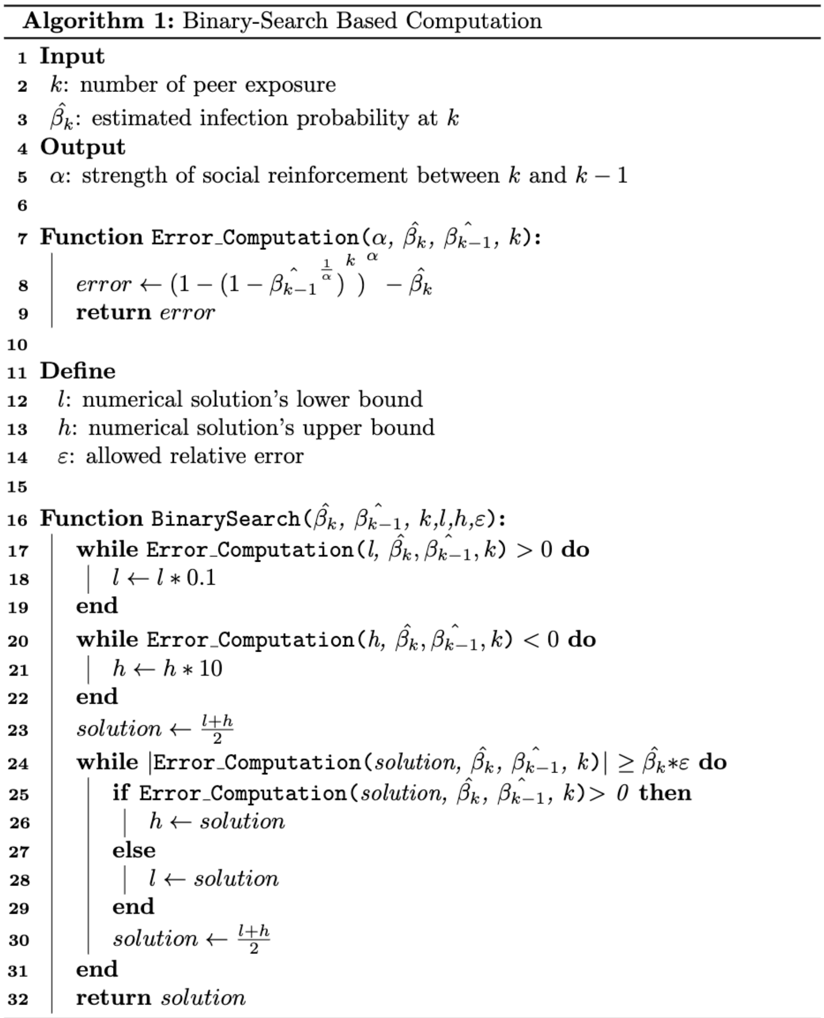

```
Algorithm 1: Binary-Search Based Computation
Input
  $k$: number of peer exposure
  $\hat{\beta}_k$: estimated infection probability at $k$
Output
  $\alpha$: strength of social reinforcement between $k$ and $k-1$

Function Error_Computation($\alpha$, $\hat{\beta}_k$, $\hat{\beta_{k-1}}$, $k$):
    $error \leftarrow (1-(1-\hat{\beta_{k-1}}^{\frac{1}{\alpha}})^{k})^{\alpha} - \hat{\beta}_k$
    return $error$

Define
  $l$: numerical solution's lower bound
  $h$: numerical solution's upper bound
  $\varepsilon$: allowed relative error

Function BinarySearch($\hat{\beta}_k$, $\hat{\beta_{k-1}}$, $k$,$l$,$h$,$\varepsilon$):
    while Error_Computation($l$, $\hat{\beta}_k$, $\hat{\beta_{k-1}}$, $k$) $> 0$ do
        $l \leftarrow l * 0.1$
    end
    while Error_Computation($h$, $\hat{\beta}_k$, $\hat{\beta_{k-1}}$, $k$) $< 0$ do
        $h \leftarrow h * 10$
    end
    $solution \leftarrow \frac{l+h}{2}$
    while |Error_Computation($solution$, $\hat{\beta}_k$, $\hat{\beta_{k-1}}$, $k$)| $\geq \hat{\beta}_k * \varepsilon$ do
        if Error_Computation($solution$, $\hat{\beta}_k$, $\hat{\beta_{k-1}}$, $k$)$> 0$ then
            $h \leftarrow solution$
        else
            $l \leftarrow solution$
        end
        $solution \leftarrow \frac{l+h}{2}$
    end
    return $solution$
```

Algorithm A1. Pseudo-code for a binary search to solve Eq. 1 and find α.

REFERENCES

An, X., Gabert, K., Fox, J., Green, O., & Bader, D. A. (2019). Skip the intersection: Quickly counting common neighbors on shared-memory systems. *2019 IEEE High Performance Extreme Computing Conference, HPEC 2019*.

Aral, S. (2016). Networked Experiments. In Y. Bramoullé, A. Galeotti, & B. Rogers (Eds.), *The Oxford Handbook on the Economics of Networks*. Oxford University Press.

Aral, S., Muchnik, L., & Sundararajan, A. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(51), 21544–21549.

Aral, S., & Nicolaides, C. (2017). Exercise contagion in a global social network. *Nature Communications*, *8*, 14753.

Aral, S., & Walker, D. (2014). Tie Strength, Embeddedness, and Social Influence: A Large-Scale Networked Experiment. *Management Science*, *60*(6), 1–19.

Aronow, P. M. (2012). A General Method for Detecting Interference Between Units in Randomized Experiments. *Sociological Methods and Research*, *40*(1), 3–16.

Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. (2012). The Role of Social Networks in Information Diffusion. *Proceedings of the 21st International Conference on World Wide Web*, *1201.4145*, 519–528.

Barrat, A., Barthelemy, M., & Vespignani, A. (2008). *Dynamical processes on complex networks*. Cambridge university press.

Centola, D. (2010). The spread of behavior in an online social network experiment. *Science (New York, N.Y.)*, *329*(5996), 1194–1197.

Eckles, D., Kizilcec, R. F., & Bakshy, E. (2016). Estimating peer effects in networks with peer encouragement designs. *Proceedings of the National Academy of Sciences*, *113*(27), 7316–7322.

Eckles, D., Mossel, E., Rahimian, M. A., & Sen, S. (2019). Long Ties Accelerate Noisy Threshold-based Contagions. *SSRN Electronic Journal*.

Ferrali, R., Grossman, G., Platas, M. R., & Rodden, J. (2020). It takes a village: Peer effects and externalities in technology adoption. *American Journal of Political Science*, *64*(3), 536–553.

Fink, C., Schmidt, A., Barash, V., Cameron, C., & Macy, M. (2016). Complex contagions and the diffusion of popular Twitter hashtags in Nigeria. *Social Network Analysis and Mining*, *6*(1), 1–19.

Gerber, A. S., & Green, D. P. (2012). *Field experiments: Design, analysis, and interpretation*. W.W. Norton.

Imbens, G. W., & Angrist, J. D. (1994). Identification and Estimation of Local Average Treatment Effects. *Econometrica*, *62*(2), 467–475. https://doi.org/10.2307/2951620

Iyengar, R., & Berger, J. (2014). *How the Quantity and Timing of Social Influence Impact Product Adoption*.

Karsai, M., Iñiguez, G., Kaski, K., & Kertész, J. (2014). Complex contagion process in spreading of online innovation. *Journal of the Royal Society Interface*, *11*(101), 20140694.

Mønsted, B., Sapieżyński, P., Ferrara, E., & Lehmann, S. (2017). Evidence of complex contagion of information in social media: An experiment using Twitter bots. *PLOS ONE*, *12*(9), e0184148.

Pastor-Satorras, R., & Vespignani, A. (2001). Epidemic spreading in scale-free networks. *Physical Review Letters*, *86*(14), 3200–3203.

Rubin, D. B. (1980). Randomization Analysis of Experimental Data: The Fisher Randomization Test Comment. *Journal of the American Statistical Association*, *75*, 591–593.

State, B., & Adamic, L. (2015). The diffusion of support in an online social movement: Evidence from the adoption of equal-sign profile pictures. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 1741–1750.

Taylor, S. J., & Eckles, D. (2018). Randomized experiments to detect and estimate social influence in networks. In S. Lehmann & Y. Ahn (Eds.), *Complex Spreading Phenomena in Social Systems* (pp. 1–24). Springer.

Ugander, J., Backstrom, L., Marlow, C., & Kleinberg, J. (2012). Structural diversity in social contagion. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(16), 5962–5966.

Weng, L., Menczer, F., & Ahn, Y.-Y. (2013). Virality Prediction and Community Structure in Social Networks. *Nature Scientific Reports*, *3*(2522), 1–6.