© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

CHAPTER 1

Introduction

The promise of viral diffusion is all around us. We all know that new ideas can spread with the remarkable ease of a virus. Yet we also know that social innovations that can benefit society often fail to diffuse. The topic of this book is a new approach to using the pathways of network diffusion to accelerate social change.

A good example of a situation where this approach was successful was in Korea at the start of the 1960s. At the time, population growth rates were skyrocketing. Korea was facing an imminent population explosion. To intervene, the Korean government instituted a nationwide contraceptive initiative. Similar policy initiatives were attempted during the 1960s and early 1970s by the governments of several developing nations. They faced a similar problem. Living conditions were improving, but childbearing norms in rural households, in which families typically had five or more children, were still guided by traditional concerns of early life mortality.[1]

Most interventions were based on psychological models of behavior change. In some countries, mass-media campaigns shamed families for having too many children and attempted to induce contraceptive use by emphasizing individual accountability. The modest success of many of these programs stood in stark contrast to the Korean initiative, which surpassed all of its stated policy goals in less than twenty years. The success of this program signaled that a new way of thinking about public health interventions was on the horizon—a sociological way of thinking about how peer networks could be used to change social norms.[2]

The Korean intervention presented villages throughout the country with a menu of contraceptive options. Although Korea's program was nationally focused, its effectiveness hinged on villagers getting local exposure to contraceptive choices through social contact with their neighbors. Peer-to-peer networks of social diffusion successfully reached large numbers of adopters in many of the villages. When diffusion

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

succeeded, women tended to adopt the same contraceptive methods as their contacts. This produced uniformity on contraceptive methods used within villages; however, there was a surprising amount of variation in the methods adopted across villages. Some were "IUD" villages, whereas others were "pill" villages, and still others were "vasectomy" villages. Interestingly, the particular method of contraception was not the determining factor for successful diffusion; rather, it was the network of social influence.[3] In the most successful villages, closely knit groups were linked together by overlapping social ties, which fostered the spread of contraceptive use throughout the community. The more studies that followed, the more findings supported the same basic conclusion—that social networks are the primary pathways for the spread of new social norms.[4]

An unexpected puzzle arose, however, from the fact the network pathways that were most successful for spreading behavior change were not the same networks that would be predicted by the theory of viral diffusion. While the viral model suggests that radiating networks of weak ties would lead to successful dissemination, it was instead overlapping patterns of spatial interaction that were the key to widespread adoption. In the decades since, scores of similar findings have surfaced in every field of diffusion research, from the spread of digital technologies to the mobilization of social movements. A growing catalog of studies has found that closely knit, densely overlapping networks are associated with the successful spread of innovative behaviors.

Today, the notion of virality animates the research agendas of hundreds of thousands of scientists worldwide, ranging from computer scientists and physicists, to sociologists and marketing scholars. Across many of these areas, lessons from the field of infectious-disease epidemiology provide a general orientation for studying behavioral contagions. The guiding assumption is that behaviors spread like viruses. The author of *The Tipping Point*, Malcolm Gladwell crystallized this idea: "I'm convinced that ideas and behaviors and new products move through a population very much like a disease does. This isn't just a metaphor, in other words. I'm talking about a very literal analogy. . . . Ideas can be contagious in exactly the same way that a virus is."[5]

This book offers a different perspective on diffusion. I show why the disease theory of diffusion does not work for understanding the spread of most behaviors and what this tells us about the kinds of social networks that are best suited for spreading innovations. This journey to

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

discover how behaviors spread reveals the specific features of network structure that control the diffusion of behavior and, ultimately, shows how these features can be used to influence the process of social change. While research on diffusion often focuses on how to improve the qualities of a product or idea to make it more contagious, I consider situations in which the innovation itself cannot easily be changed. Instead, I focus on how changes to the social network of a population can transform a failed technology into a successful innovation. To demonstrate the impact of these ideas, this book is dedicated to providing practical solutions to problems of diffusion. The results offer a way of thinking about the network dynamics of social change that gives new life to the promise of using online technologies to promote sustainable changes in population behavior.

The examples used in this book vary widely, ranging from the diffusion of social media technologies to the spread of prophylactic measures for HIV to the growth of rebellion in post-Revolutionary France. The majority of examples are drawn from the diffusion literatures that I have been immersed in the longest—namely, the spread of health technologies and the mobilization of social movements. While on the surface these two topics seem to have nothing in common with one another, beneath the surface they have a shared logic of social influence. From a networks perspective, the common structures that underpin diffusion in both of these settings reveal the basic network characteristics that may be useful for improving the spread of behavior in a variety of contexts.

The findings here help to identify the kinds of networks that may be effective for spreading smoking cessation, as well as the network structures that can accelerate organizational change. These results show how to create online networks that can improve the adoption of new exercise behaviors. And they also reveal the differences between using social media to diffuse contagious memes versus to mobilize political activism. Here the dynamics of both informational and behavioral diffusion are explained within a framework that allows each to be understood on its own terms. The findings suggest a way for theorists and practitioners who are interested in diffusion to gain insight into when social networks will be helpful for spreading changes in behavior and how to make practical use of them.

One point worth stressing at the outset is that the approach here differs from approaches to social change that are based on the assumption

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

that people's choices can be altered by exposure to the right kinds of messages. This is true in many circumstances. But the present approach is *collective* rather than *individual*. One surprisingly helpful way of thinking about this is by analogy with schooling among fish. Studying fish individually, it would be impossible to anticipate the complex schooling behaviors that they produce when they interact as a group. Similarly, studying people one at a time provides little insight into the collective dynamics by which new behaviors spread through a population. Diffusion, like schooling, is a collective social process that unfolds through the complex interactions of many interdependent actors. The approach adopted here is to study behavior change as we would study schooling—not as an individual phenomenon, but as a collective one. This perspective assumes that people are often in situations where the decisions they make are influenced less by the information they have access to, and more by the social norms that are common in their networks. The goal here is to show how these social networks may themselves be used to control the schooling process, and spread lasting changes in behavior.

ISN'T IT OBVIOUS?

Science has often been described as the development of new intuitions about how the world works. Commentary on the science of sociology has noted that while much of contemporary sociology can seem obvious today, it was not always so. Ideas that may seem bromidic now were once revolutionary approaches to thinking about social problems. The seemingly inevitable fate of successful ideas is to be absorbed into the body of scientific knowledge, eventually entering the popular lexicon, where they are reduced from novel intuitions to tacit features of everyday life. However, there are also scientific ideas that are so counterintuitive that they defy integration into the body of popular knowledge. These intuitions present such a challenging contrast with the expectations forged by a long evolutionary, cultural, and personal history that they are hard to hold on to even once they have been learned.

A quick example here will illustrate what is meant by a counterintuitive idea and how it can happen that a scientific discovery can remain counterintuitive even once it has been explained. Figure 1.1 shows a picture of two coffee tables. The intuition that I want to elicit concerns which of the two tables is longer. Look at each table and consider the ratio of its length to its width. What would you say it is? When

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

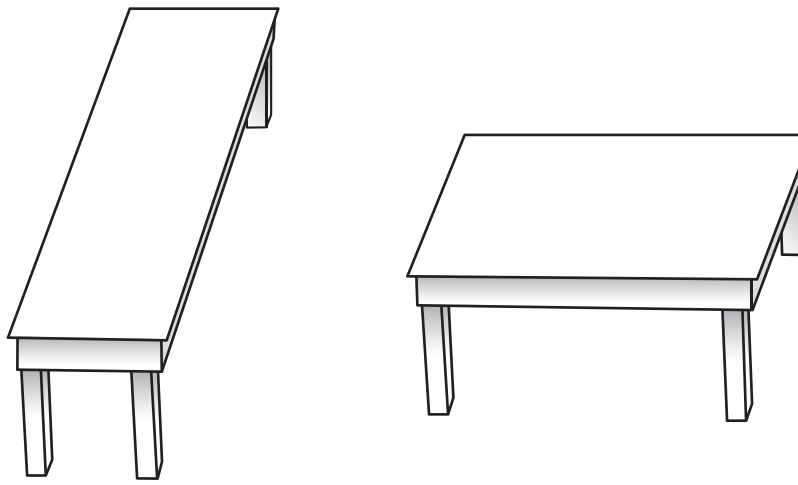

Figure 1.1 Adapted from Richard Thaler and Cass Sunstein, *Nudge: Improving Decisions About Health, Wealth, and Happiness* (New Haven, CT: Yale University Press, 2008).

I first saw this figure in the 2008 book by Richard Thaler and Cass Sunstein,[6] I guessed that the one on the left is perhaps 3:1 or 3.5:1, while the one on the right is closer to 1.5:1 or 1.25:1. Make your guess.

Now, take out your pen and lay it against the page. They are, in fact, the same table. Cognitive psychologists explain this illusion in terms of the way that the eye corrects (or fails to correct, depending on how you see it) for the orientation of the figures and the visual contrast created by the legs. Once you have measured the figures to your satisfaction and have internalized this new piece of knowledge, look away and then look back. Which table is longer?

The point is that despite having the right answer in mind, the objects nevertheless look the same as they did before. The bias in the perceptual system cannot be overcome by the knowledge that it is there. The value of scientific education is that once the bias is explained, a person can anticipate this kind of error and take precautions to avoid making mistakes in situations where it might matter. Whenever vigilance is surrendered, however, even if for a moment, a particularly persistent illusion can lead the mind to make unavoidable, and quite consequential errors in judgment.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

This book is about just such an illusion, but not one in the perceptual science of psychology. Rather, it is about a similar kind of bias in our understanding of social networks. In particular, it is about a compellingly intuitive theory of diffusion that, like the apparent differences between the two tables in figure 1.1, is likely to be persistent. Nevertheless, the intuitive appeal of this idea notwithstanding, this book shows how this popular and intuitive theory of diffusion can go seriously wrong, leading to costly errors in our understanding of how behaviors spread through social networks. The intuitive theory I am talking about is called the *strength of weak ties*.

OUTLINE OF THE CHAPTERS

The basic idea of the strength of weak ties is that while our *strong ties*—that is, our friends and close family—all tend to know each other, our *weak ties*—that is, our casual acquaintances –connect us to remote parts of the social network. As the sociologist Mark Granovetter famously put it, "Whatever is to be diffused can reach a larger number of people, and traverse a greater social distance, when passed through weak ties rather than strong."[7] Our journey here starts in chapter 2 with the initial finding that launched my work into this topic—namely, that there is an unexpected problem with this remarkably influential theory of network diffusion.

The broad influence of this theory is due in part to the recent explosion of network science across disciplines such as physics, biology, and computer science, which ushered in a period of rapid discovery for understanding how the structure of social networks affects the dynamics of diffusion. What all of these fields have in common is a belief in the idea that a contagion, such as a virus, an idea, a meme, a method of contraception, a diet, a fashion, an emotion, an ideology, or a technology, can spread from one person to another. The guiding principle of all of this work is that the structure of social contacts can foretell how a contagion will diffuse through a population. The full impact of Granovetter's original insight was not realized until the physicists Duncan Watts and Steven Strogatz developed the *small-world* model, which demonstrated that bridge ties—that is, social links connecting otherwise distant people—can dramatically increase the rate of diffusion across social networks.[8] The strength of weak ties hypothesis and the small-world principle resonate with one another to present a unified and powerful view of how network structure controls the dynamics of

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

social diffusion. The problem is that when we compare this view to a large body of empirical research on diffusion, a puzzle arises from the fact that while weak ties seem to improve diffusion in some cases, there are many other cases in which they do not.

The solution to this puzzle comes in chapter 3, with the finding that there is an important difference between "complex" behavioral contagions, for which transmission requires contact with multiple adopters, and "simple" informational and viral contagions, for which transmission only requires contact with a single source. Computational explorations show that when contagions are complex, because they are costly, risky, or involve some degree of complementarity, weak ties can slow down diffusion. This finding has implications for most of the contagions that social scientists care about, such as cooperation, social norms, marriage practices, health behaviors, voting behavior, technology adoption, and investment decisions, to name just a few.[9] It also means that social networks that accelerate the spread of an infectious disease can slow down the diffusion of its cure. This occurs because diseases, like information, are typically simple contagions that pass quickly along weak ties. Behavior change, however, typically is not.

With this finding, chapter 4 turns our attention from the mathematical world of computational experiments to the empirical world of behaviors spreading through human social networks. This is where we face a crucial challenge—devising a way to test this theory of diffusion empirically. For the vast majority of research on networks and diffusion, even the rudimentary task of identifying the existence of a diffusion process has been fraught with difficulties, to say nothing of being able to identify exactly how the structure of a social network may have altered it. Here the Internet is an invaluable ally for social research. Over the course of two years, an independent online community was constructed and populated with thousands of volunteers recruited at large from the World Wide Web. Techniques from small-group laboratory experiments were combined with tools from large-scale data science analytics to conduct an Internet-based social network experiment of how behaviors spread through online communities. The illuminating results from this study show that while weak ties were highly effective for spreading information, they slowed down the spread of behavior.

These results suggest that the rapid diffusion of information through weak ties may not tell much about the dynamics of behavior change. In fact, the more quickly that information goes viral, the less promising

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

the outlook may be for spreading behavior. Thus, the finding that emerges from the intuitive distinction between simple viral contagions and complex behavioral contagions is the counterintuitive insight that the more weak ties there are in a network, the slower that innovations may spread.

In part 2 of this book, I use this theory of social contagions to address practical problems of diffusion. Chapter 5 shows the range of empirical settings to which the theory of complex contagions has been applied—from the spread of political hashtags on Twitter to the diffusion of smoking among teens.

Chapter 6 shows how these findings can be used to address the specific challenges that arise when innovators face social opposition. One application shows how public health interventions may be designed in order to trigger network cascades of behavior change in at-risk populations. Another application considers how social networks can be used to incubate the spread of an innovative technology in a population where an alternative product is already entrenched. In each case, the lesson is the same: clustering the early adopters together can increase the spread of innovation.

Chapter 7 turns to the topic of organizational performance and shows how the findings in this book challenge conventional wisdom about the value of information brokers for diffusing innovations. This chapter identifies the importance of *wide bridges* for spreading new behaviors and ideas across organizational boundaries. The discussion here also explores the origins of network structure. This chapter shows how the identities that people have within an organization can influence the structure of the networks that emerge, and demonstrates how organizational identities can be used to design networks that are effective for diffusion.

Building on these practical applications, part 3 takes a hands-on approach to constructing new forms of social capital online. Chapter 8 offers experimental findings on how to design social networks among strangers to increase the flow of new behaviors. The results highlight the importance of both social relevance and empathy in network ties and show how these factors can be strengthened within existing online settings by incorporating *homophily*—that is, similarity between social contacts—into the architecture of a social network.

Chapter 9 then turns to the difficult problem of how to control the kinds of behaviors that spread online. Social influence comes in all shapes and sizes, and there are some circumstances in which

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

constructing influential networks may backfire by spreading undesirable behaviors. The relational context of social networks comes to the foreground here. The results show that sometimes the most intuitive network strategies for inducing behavior change can have the least desirable outcomes. To offer some guidance on how to avoid this, chapter 9 identifies how features of social comparison and social support in online network settings can determine the kinds of influences that people will have on each other's behavior. A policy experiment illustrates these ideas by showing how the design of relationships within an online community can catalyze, or inhibit, changes in physical activity.

By the end of this book, the discussion has developed from studying the effects of strong and weak ties on diffusion to demonstrating how the principle of social reinforcement gives new insight into the network dynamics of behavior change. The basic approach throughout is always the same: seeing how imperceptible changes in the structure of social relationships produce significant differences in collective outcomes. This method allows more than the understanding of individual behavior: it provides an appreciation of the unseen forces that guide the movements of collective behavior. The most promising finding is that the reasonable expectation that people will resist behavior change does not mean that people are incorrigible. Nor does it mean that diffusion will fail. Instead, this expectation reveals the pathways that behavioral contagions will need to follow if they are to flow through a population—and the strategies that can be used to make this process most effective.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

For general queries, contact webmaster@press.princeton.edu