

Creating Persuasive Technologies: An Eight-Step Design Process

BJ Fogg
Persuasive Technology Lab
Stanford University
captology.stanford.edu
www.bjfogg.com
bjfogg@stanford.edu

ABSTRACT

This paper outlines eight steps to follow as best practices in the early stages of persuasive technology design. The eight-step process, drawn from demonstrated successes in industry practice, begins with defining the persuasion goal to match a target audience with an appropriate technology channel. Subsequent steps include imitating successful examples of persuasive design, performing rapid trials, measuring behavioral outcomes, and building on small successes.

General Terms

Design, Measurement, Experimentation, Human Factors.

Keywords

Persuasion, design, persuasive technology, prototyping, iteration, behavior change, captology, behavior model.

INTRODUCTION

Fifteen years ago, there were relatively few examples of persuasive technologies in our lives. The web wasn't ubiquitous, and software wasn't designed to change behaviors; it was focused more on crunching data and boosting productivity. But today persuasive technologies are ubiquitous; we are surrounded by digital products designed to change what we think and do. Persuasive technology experiences come to us through the web (from commerce sites to social networking), video games (e.g., Wii Fit and Dance Dance Revolution), mobile phones (e.g., health applications for iPhone and commercial texting services), and specialized consumer electronic device, from "talking" pedometers to bathroom scales that track body mass.

Increasingly, the living room TV and even automobiles are channels for persuasive experiences. For instance, TiVo not only suggests programs to watch but integrates Netflix and encourages

customers to make purchases on Amazon. As for automobiles, one feature of the Toyota Prius is a miles-per-gallon meter that motivates owners to adopt more eco-friendly driving habits.

Today those of us who are interested in the design and study of persuasive technologies have a wealth of examples from which to choose. The existence of so many successful examples changes the study of persuasive technology in significant ways. We no longer have to invent new persuasive solutions out of whole cloth. Instead, we can focus on existing persuasive technology products and techniques, varying those systems to understand the dynamics and principles of persuasive design. In this way, we can learn most rapidly about the psychology of persuasion and persuasive technology by working with existing solutions.

That said, there still will be times, either for commercial purposes or for our own academic research, when we want to create an entirely new persuasive technology for which there is no good prototype. This can be a challenge, given that many people have little or no experience in creating products with a persuasive goal, and our emerging field does not yet have a systematic design process. The lack of a well-defined process for designing persuasive technology leads people to adapt methods from other fields, such as usability engineering, or to make guesses as how to define and develop their products. Neither approach is efficient.

To address this challenge, in this paper I draw on my 15 years of experience in studying and creating persuasive technologies to offer what I consider to be "best practices" for developing new digital experiences that influence people. I share and explain an eight-step process for creating successful persuasive technologies, focused on early-stage design. These steps are based on a combination of my academic research, the many project-based courses I've taught at Stanford University, and my work in industry designing and testing persuasive technology solutions. My goal in writing this paper is to help academics and practitioners get started down the most efficient and promising path for creating persuasive technologies.

THE PROBLEM: BIG FAILURES

As I've reviewed a broad range of work from students and companies, I've found that attempts to create persuasive technologies often fail. One problem is that many projects are too ambitious, and thus are set up for failure. For instance, a design team might select a challenging behavior, such as smoking

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Persuasive '09, April 26-29, Claremont, California, USA.

Copyright © 2009 ACM ISBN 978-1-60558-376-1/09/04... \$5.00

cessation, as the target. While everyone would agree that helping people to quit smoking is a good idea, such a target is too ambitious for design teams that have never before created a persuasive technology. Stopping smoking—especially for those who have a long-time habit—is the Mount Everest of human behavior change. Just as no one would expect a novice to succeed in reaching the summit, novice design teams should scale back their ambitions and save the difficult behaviors, such as smoking or weight loss, for later projects, after they have learned to succeed in designing technologies targeted at more tractable behavior changes. As with any other endeavor, acquiring skill in creating persuasive technology takes practice. Most attempts will fail, but experience increases one’s ability to create successful products.

INCREASING THE ODDS OF SUCCESS

The purpose of the eight steps I propose in this paper is to outline a path to follow in designing persuasive technologies that will increase the probability of success. The process starts with careful thinking, then introduces small, simple tests to produce measurable success. Once a design team finds success, albeit small, only then should the team attempt to achieve larger, more ambitious goals. This is the route that leads to growing success rather than ongoing frustration.

In recent years, I’ve encouraged academics and industry players to think small while they are starting their design work in persuasive technology. I’ve advised people to be less ambitious on project goals and focus first on achieving small successes.

The reaction I get is sometimes negative. This doesn’t surprise me. Many people are attracted to persuasive design because they want to do something big and important with technology. In addition, at times people feel compelled by an outside force—a grant, for instance, or pressure from a company leader—to set out on the most ambitious course. In the face of long odds, novice design teams may insist on tackling the hardest issues, such as alcoholism, weight loss, smoking, or global warming, because of such outside pressures. These attempts almost always fail, either because the project never gets completed, or tests of the technology show that it has no lasting effects on behavior.

Most of the persuasive systems we see at conferences or read about online are the few projects that succeeded. In this context, where only the successes are visible, it’s easy to forget that the landscape of persuasive technology is riddled with the carcasses of failed projects. In fact, I’m certain the failures vastly outnumber the successes.

Failure isn’t always bad. In fact, best practice in persuasive design includes expecting to fail. But I advocate *failing fast*. This means investing only a few hours in a trial that might not work. In contrast, projects that require months or years should not fail; that’s a waste of time and money. Large projects will succeed when built on a foundation of many small, measurable successes.

Learning how to succeed with persuasive design projects in the early stages is the purpose of the steps I propose in this paper.

Before I go on, a few words about my bias: While I view persuasion as either attitude or behavior change, in this paper I focus on *behavior change* in persuasive technology. I believe that changing behavior is what matters most in issues of politics, health, environment and more. Fortunately, measuring behavior change is getting easier with today’s technology.

THE EIGHT-STEP DESIGN PROCESS

The eight steps in the process of designing persuasive technology, described below and outlined in Figure 1, are carried out mostly in sequence. In some cases, two steps may be carried out in parallel; at other times, the design team may back up a step and re-think or re-try. The eight steps are not intended to be a rigid formula; instead, the steps serve as milestones to make the design process more effective. Varying the sequence of steps to suit the circumstances is a valid part of the design process.

Step 1: Choose a simple behavior to target

The first step in designing a successful persuasive technology is to select an appropriate behavior to target for change. The design team should select the smallest, simplest behavior that matters. Often this requires a team to reduce their big goal to a small, seemingly tiny, objective.

For example, I worked with a large health care company whose goal was to help people reduce their stress levels. As an objective for a persuasive technology design, that goal is too vague and ambitious. So we sought to reduce the big goal to something more useful for our early design purposes. After brainstorming many options, we decided to target the following desired behavior: Stretch for 20 seconds each day when prompted.

Note that this smaller goal was so simple that anyone could achieve it, and success was measurable. Of course, we realized that getting people to stretch for 20 seconds wasn’t our final objective. But this simple behavior was a good starting point for the larger goal: reducing people’s overall stress level.

A large, vague goal can be broken down in two ways. Sometimes a small goal can be an approximation of a larger objective; for instance, stretching is a small solution for reducing stress. Other times, a small goal can be a first step in achieving the larger goal. For example, watching a short video online about Pap tests can be the first step in persuading women at risk to actually get tested.

Achieving the small goal may have bigger effects than expected; persuasion professionals have long understood that getting people to do small things naturally leads to their adopting more ambitious behaviors, even without a bigger intervention. For example, if the big goal is to get people to be more environmentally friendly, the small step of motivating them to change a single light bulb in their home could alter how they view themselves, and over time they are more likely to make other eco-friendly choices that are consonant with this small step.

This first step in the process—choosing a simple behavior to target—is the most important aspect of designing successful persuasive technologies. But taking this step is more difficult than it sounds. The best designers will advocate for simplicity, but doing so takes courage. When working with a team or under supervision, a designer risks looking timid when saying “no” to an ambitious goal and proposing instead something small and simple. In some situations, members of the design team may already have their own pet ideas for what they want to build, so the team collectively nods while each person adds an additional bell or whistle to the user experience, complicating the project and unwittingly setting the team up for failure.

Step 2: Choose a receptive audience

Step 2 in the persuasive design process involves choosing the right audience for your intervention. Where designers have a choice (i.e., the audience is not pre-determined by the project), I advocate

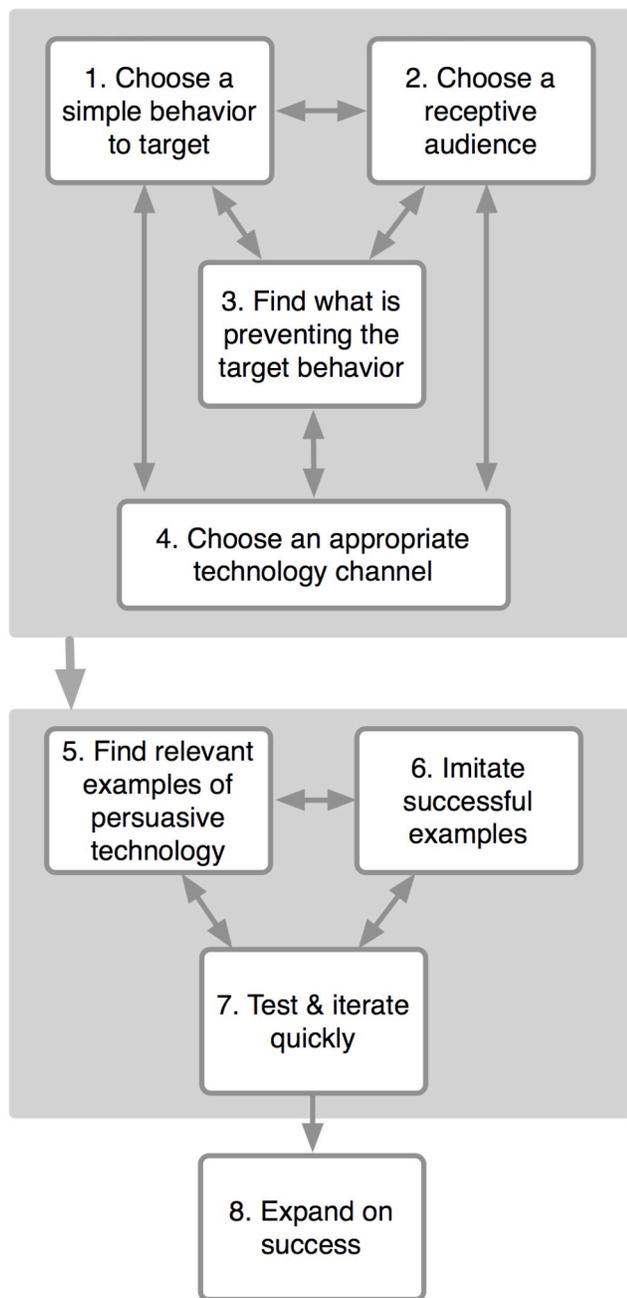


Figure 1: Eight steps in early-stage persuasive design

choosing the audience that is most likely to be receptive to the targeted behavior change. The audience also should be familiar with the technology channel (I will discuss channel in more detail in Step 4).

The team may be tempted to design an intervention for the toughest audience, such as helping compulsive gamblers to stop. In my experience, this is a mistake. A related mistake is to design the intervention for all users rather than a specific user type. Neither approach works well.

The goal of the Steps 1 through 7 of the design process is to create a digital product that reliably persuades someone—not *everyone*—to adopt the target behavior. In Step 8, I discuss expanding the audience, but this step should be carried out only after the

technology has been successfully tested on a more responsive audience. Until then, I advocate choosing the easiest target audience. For example, if the team is designing a technology to persuade users to adopt better eating habits, they should select an audience who has demonstrated a desire to improve their diets. If a team wants to persuade people to adopt a daily exercise routine, designers will increase their odds of success by focusing first on people who already exercise once in a while.

Design teams have so many things to worry about when creating a new persuasive technology that a resistant audience is not helpful. In fact, choosing the wrong audience will almost certainly doom the design project, especially in the early stages. As I will discuss in later steps, once a design team has developed an intervention that is working, they will be able to expand their target audience and bring in users who are less receptive to the intervention.

The next consideration in choosing an audience is how familiar people are with technology. I advocate choosing early adopters or other adventurous souls as a target audience. I believe it's a mistake to target an audience that is afraid of computers or is just beginning to use the technology channel for which the team is building a persuasive technology, be it texting, social networking, or interactive TV. The best audience for early projects consists of those who enjoy using technology and trying new things.

In some cases, the first two steps of the design process might be completed in reverse order. Sometimes the audience will determine the target behavior, rather than vice versa. For example, a project to motivate teens to save money is likely to target a different behavior (e.g., getting into the habit of saving) than a project to persuade older adults to save (setting aside a specific amount to ensure a secure retirement). So if work in Step 2 causes a team to back up to Step 1, that's okay; finding the right combination of behavior and audience is vital to laying the foundation for the subsequent steps in the design process.

Step 3: Find what prevents the target behavior

Once a design team has selected the appropriate behavior and audience to target, it's time to move on to Step 3. In this step the team must determine what is preventing the audience from performing the target behavior. For example, if children in first grade aren't brushing their teeth each morning, what is lacking? As another example, if alumni aren't donating to their alma mater, why not?

The answers to such questions always fall into some combination of the following three categories:

- lack of motivation
- lack of ability
- lack of a well-timed trigger to perform the behavior

In other words, in Step 3 the design team must pinpoint why people aren't performing the behavior. Is it because they are not motivated to perform the behavior? Is it because they lack ability? Or is it because they are not being triggered to perform the behavior at the right time? Or is it a combination of the three factors? The answers in Step 3 will determine the work required in later steps, so a thorough examination at this stage is critical. (For more details about this method, see www.BehaviorModel.org)

Consider, for example, a middle-class family living in the suburbs of Los Angeles. Suppose that the family is not at all eco-friendly. If the design team's goal is to motivate the family to use eco-

friendly light bulbs, in Step 3 the team should explore why the family is not already doing so.

The design team may find that the family has both the motivation and ability to use eco-friendly bulbs, but they are lacking a trigger to perform the behavior. Or, the team may discover that the family lacks motivation: They don't see the benefits of using eco-friendly bulbs. Or perhaps the family lacks ability—they don't know which bulbs to purchase, or they can't afford them, or they feel they don't have time to change out the traditional bulbs in their homes.

Technology interventions that require only a trigger are the easiest to create and the most likely to succeed. For example, in the stress reduction project described earlier, in which the "small goal" was to persuade people to stretch for 20 seconds, participants in our pilot needed only to be reminded (one type of trigger) to stretch. We didn't need to motivate them to stretch, or teach them how.

In practice, a persuasive technology solution will often require more than simply triggering a desired behavior. Rather, the solution must also boost motivation or facilitate the behavior, or both. If the target audience lacks only motivation, the persuasive design should focus on motivation. If ability is lacking, the solution should facilitate the target behavior. One caution: If the target audience is lacking *both* motivation and ability, the team may want to back up and rethink the previous steps. Early-stage teams usually struggle with the task of creating a persuasive technology that simultaneously motivates and facilitates a target behavior. Achieving both at the same time is hard. In such cases, I suggest redefining the target behavior or the audience.

Step 4: Choose a familiar technology channel

Once a design team has identified what is preventing people from adopting the target behavior, they can move on to Step 4: choosing the best channel for the technology intervention. Which channel is "best" usually depends on three factors: the target behavior, the audience, and what is preventing the audience from adopting the behavior—i.e., the first three steps in the design process. What this means is that in most cases, the design team cannot select an intervention channel—web, mobile phone, video game, or other—until the first three pieces of the process have been completed. (I'll address exceptions to this in a moment.)

Today, we have an increasing number of technology channels for persuasion: Web, software installed on personal computers, mobile phone applications, texting on mobile phones, social platforms like Facebook, online video, platform games, and so on. The challenge for the design team is to choose among the range of persuasion channels available, considering how well each channel matches the target behavior. For example, if a design team is creating an experience to motivate donations to a political party, the team will likely need to use the web in their solution to enable the financial transaction. If the target behavior is to share a message with at least one friend, then the channel could be email, online video, or social networks, because all of those channels make sharing easy.

The next issue in channel selection is audience. A design team must select a channel that is familiar to the target user. I've watched teams expect their audience to learn a new channel (such as texting or social networking) and simultaneously adopt a new behavior. This approach almost never works. I have come to believe that most people can change only one behavior at a time. And the reality is that adopting a new technology is a behavior change. It's unrealistic for designers to think they can layer in

another behavior change, such as daily exercise, without overwhelming their audience.

If a design team must use a channel that is unfamiliar to the audience, there is a process for doing so, but it takes time. To use a new channel to change behavior, design teams must first help the audience to become familiar with the channel. For example, if AARP wanted to motivate older adults to walk each morning, using text messaging as the channel for triggering the behavior, the AARP design team would first need to educate their audience in how to use texting, which in the U.S. is not common for older adults. Only after their audience was comfortable with texting should the team introduce an intervention to promote walking.

Training people to use a new channel is difficult, which is why I suggest avoiding it when possible. The easiest and fastest way to progress with persuasive design projects is to select a channel the audience already uses. In some cases, this may limit the intervention to widely used channels, such as email, web sites, and online video. In contrast, mobile texting, video games, and social networking are not familiar to everyone today and may work well only for particular types of audiences.

Finally, the design team must select a channel that addresses the answer to the questions from Step 3: Why isn't the audience performing the target behavior? Is the problem a lack of motivation, ability, triggering, or some combination? The answer will help to guide channel selection. Some channels, including online video, social networks, and video games, are effective at increasing motivation. Other channels, such as installed software and specialized devices, excel at making a behavior simpler (which is functionally the same as increasing ability). And some channels, such as text messaging and email, work well for triggering behaviors.

As an example, consider again the family in Los Angeles that is not behaving in eco-friendly ways. If the family is primarily lacking motivation, the design team should consider channels that leverage motivation, such as social networks and online video. If ability is missing, the team could consider a web service that makes the behavior easier, such as a guide showing where to buy eco-friendly light bulbs. Online video also could be used to enhance the target audience's ability to perform behaviors, such as a video that shows the process for replacing light bulbs, step by step. If the family is lacking only a trigger to change their light bulbs to more eco-friendly versions, then email or text messaging may be the simplest solution.

Re-ordering the First Four Steps

Usually, the first four steps in the persuasive technology design process are performed in sequence. But in some cases a design team will make an exception and carry out the steps in a different order. For example, if the design team works for a health insurance company, the team might be assigned an audience for intervention, such as sedentary older adults; the team may not have a choice. But they might have lots of flexibility about the channel. In contrast, a start-up company that provides mobile services will not have flexibility in terms of the channel choice (they will use mobile, of course), but they may be entirely flexible on target behavior and audience. Finally, consider the case of a graduate student designing persuasive technology. The student may not be personally attached to a particular behavior, audience, or channel. But the student might have a constraint, such as a funding source that determines the behavior, an advisor who wants work done with a specific audience, or a class project that focuses on a channel, such as social networks.

Designers work from different starting points, with different constraints. As a result, they may carry out the first steps in the design process in a different order. As noted earlier, this eight-step process is a guideline, not a rigid mandate. Designers should apply it to the unique circumstances of each project. That said, whichever sequence the design team follows, the first four steps should come before moving on to Step 5.

Step 5: Find relevant examples of persuasive technology

In Step 5 of the design process, the team should search for examples of successful persuasive technologies that are relevant to their intervention, as defined in the previous steps. Suppose the behavior of interest is persuading people to donate money to a particular cause. The design team must find examples of persuasive technology solutions that succeed in getting people to make donations. If older adults are the target audience, the team needs to find existing solutions that work for that audience. Or if the chosen channel is video games, the team needs to study video games that have successfully changed behaviors.

A design team won't always know if a given persuasive technology is successful, because companies generally don't share their conversion data with outsiders. So in Step 5, making educated guesses is a good approach. For example, if a leading Web 2.0 company is using reminder emails to trigger behavior—for instance, driving customers to the company's website—those emails are likely persuading effectively, or the company would probably not continue this approach. Similarly, if a leading social networking site attempts to get people to invite more friends using a certain widget, then that widget is likely achieving the company's objective. In other words, a design team should find relevant examples from companies that are succeeding, because their solutions for persuasive design, even on small issues, are likely to be effective.

In searching for relevant examples of successful persuasive technology, design teams will rarely find one example that matches the precise behavior, audience, and channel the design team has chosen for its project. And even if this parallel example were to exist, the team would want to study other solutions to see a range of options. Specifically, a design team should examine at least nine examples in total: three that achieve a similar behavior, three that reach a similar audience, and three that use the same technology channel as the design team's.

Teams can also learn to design for persuasion by imitating the methods of experts who work in that domain. For example, if the team's intervention focuses on persuading donations, a team can learn from studying the best practices of professionals who persuade people to donate.

With relevant examples of persuasion in hand, a design team is ready to move to Step 6.

Step 6: Imitate successful examples

The next step in the persuasive design process is to imitate what's working in the successful examples gathered in Step 5. Ten years ago, persuasive technology was so new that design teams needed to create novel solutions. Today, the landscape is different. Rather than starting from scratch, a better, more reliable method is to imitate what's already working—on Facebook, Amazon, in video games, and more—and adapt those successful approaches to the target behavior and audience. With so many examples of successful persuasive technologies, there is no need to reinvent the

wheel. Identifying and adapting successful technology examples to the design project at hand is the fastest, surest way to create effective persuasive technologies.

The design team should not be afraid of doing something that is similar to what has already worked. In the later stages of the design process (Step 8), the team will have many ways to be unique. The opportunity for real innovation comes after laying a solid foundation.

Besides selecting the right behavior, imitating successful examples is perhaps the most important step thus far in creating persuasive technologies. Nevertheless, I've found that teams sometimes resist this step because the result feels derivative. Innovators often want to be creative; they hope to craft something completely new. That's a mistake—at least if the goal is efficiency in creating a persuasive technology that works. When learning new skills, people succeed fastest by imitating success, and the design of persuasive technologies is no exception.

Step 6 requires insight. When a design team examines a successful example, the team must be able to identify the "secret sauce" – the special ingredient that makes the example effective. The secret sauce usually is not superficial design elements such as color or typeface. So the team's challenge is to evaluate the example from a psychological perspective to discover the essence of its persuasive power. Step 6 is easier when team members have a background in psychology and good intuition about persuasion.

The questions from Step 3 are a good starting point for Step 6: What is the successful example doing to achieve behavior change? Is it motivating? Providing ability? Triggering the audience to adopt a behavior? Often, an example will seem to address all three questions, but the team must look deeper to find the design element's primary purpose. For instance, if the persuasive element serves to boost motivation, then the team can adapt this secret sauce to its own design project.

In Step 6, the design team can use their creative skills to come up with ways to adapt the secret sauce of the existing example to their own target behavior. The natural inclination is to find one successful example, adapt it, and then stop, because it seems the problem is solved. But in Step 7 the team will need many things to test, so having just one imitation is not sufficient.

Step 7: Test and iterate quickly

After a design team has found ways to imitate successful examples of persuasive technology, the next step is to test various persuasive experiences quickly and repeatedly. A series of small, rapid tests will teach more than one big test. Each test should take only a few hours, start to finish. These are not scientific experiments but quick trials that allow the design team to prototype the experience and see how people react. The team should assess the response, ideally by measuring behavior.

For example, suppose a team wants their target audience to share an online slideshow with at least two friends. In Step 5, the team has found examples of how online slideshows are passed from one friend to the other. In Step 6, they've identified the persuasive power in these examples. Also as part of Step 6, the team has created a few variations of the successes. In Step 7, the team quickly tests one of these variations, ideally doing the easiest and fastest solution first. In this case, the fastest test might be creating an email that asks people to watch a slideshow online and then share it with friends. The team's email would imitate – in tone, length, persuasion strategy, and format – the example from Step 6. Setting up this test might take two hours. Putting the test into

motion might take another two hours. Then the team can review the results and prepare for a new test.

The tests in Step 7 are not scientific experiments to gather publishable data. Instead, the focus is on rapid trials to learn quickly about designing for persuasion, given the team's target behavior, audience, and channel.

Designing for persuasion is harder than designing for usability. Many attempts to change people's behaviors fail. That's why Step 7 calls for rapid testing and many trials. The faster a team tests various options, they faster they learn what will work. I believe it's a mistake to invest more than 10 or 15 hours on an early test. Ideally, the cycle should be just a few hours.

Furthermore, the design team should set low expectations for their trials so they don't get discouraged, and so other people who are watching (advisors or bosses, for example) won't expect success early. To manage expectations, the design team might tell themselves and others that they plan on doing ten rapid trials, learning along the way. With each test, successful or not, the team will gain more insight into rapid testing and what's likely to work the next time.

I've found that success in conducting trials results not from the size or complexity of the test; instead, success seems to correlate with the number of tests performed. As I see it, knowing how to prototype, test, and evaluate results quickly is the most valuable skill for designers of persuasive technology.

At this point in the process, design teams may go back and forth between Step 7 and Step 6. The goal is to keep running quick tests and to continue learning. Again, although results are measured, the trials should not be viewed as true experiments with protocols and statistical significance.

Like other components of the design process, Step 7 takes courage. The traditional academic approach is to spend months on a big study, planning it, running it, and analyzing the data. Similarly, with the old method technology innovation, projects took years to plan, code, debug, and distribute. Today, we have tools to create prototypes and products much faster. So even though academics and companies may cling their "slow and careful" roots, a designer of persuasive technology products benefits from quick, lightweight testing to find success.

Every successful online service I've examined in the last few years has become successful through starting small and iterating quickly. The winners in today's world of consumer Internet did not plan for years or run long trials; likewise, in the world of persuasive technology, that is a losing approach. As I have stated elsewhere, when it comes to creating consumer services, "many crummy trials beats deep thinking."

Again, the goal of Step 7 is to find something that works, to create an intervention that succeeds in helping the target audience to adopt a very simple target behavior that can be measured. Once the design team has achieved success on a small scale, they can move on to Step 8.

Step 8: Expand on success

Creating a persuasive technology that changes a behavior, no matter how small or simple, is a milestone. In Step 8, the design team can expand on this success. Now is the time to scale up.

There are a variety of ways in which the team can scale. One way is to make the target behavior more difficult. For instance, instead

of having people install one single eco-friendly light bulb into their home, the expanded intervention can focus on persuading people to replace all the inefficient light bulbs in their home. Another way to expand is to reach out to a new audience—new types of users who are less adoptive, users who are tougher cases, and see how the intervention works with this new audience. A third way is to expand the scope of distribution, reaching a wider audience with the intervention.

The method of expansion should be systematic, varying only one or two attributes from the success achieved in Step 7. For example, consider the intervention with my health-care client to persuade people to stretch for 20 seconds on cue. The compliance rate was over 50%. At that point, our team had a process that worked well for the audience. With that success, it was easy to see ways to expand. Many options existed: We could persuade the same audience to stretch for a longer time period. We could use the same method to motivate better eating habits. Or we could reach out to a more resistant audience and determine if the intervention would be successful with them.

The decision of how to expand depends on the company's goals or the researcher's agenda. In the case of our stretching intervention, we expanded to a longer, more intensive program to reduce stress, with behaviors going beyond stretching. Because we expanded the behavior, we kept other parts of our program the same: similar audience, same channel, same underlying psychology, and the same types of metrics. Our expansion worked at least as well as the original. And now another expansion is in order, one that will focus on a wider range of people who are not as familiar with texting, the channel we originally used.

From an academic perspective, Step 8—not earlier in the design process—is the starting point for a controlled experiment. The first 7 steps are focused on getting the design team to a place where success is likely to be achieved. By the end of Step 7, the pilot testing has been completed, and the research team can feel confident about the intervention process, the participants, and the independent and dependent variables. From a scientific standpoint, the true experiment doesn't begin until Step 8.

EVERYTHING BIG STARTED SMALL

When one looks at the history of successful consumer Internet services, a striking similarity emerges: Each service started in a small, focused way: Google offered a simple search box. Yahoo was merely a list of links. Facebook was a directory created for friends. As the small offerings succeeded, they expanded. That approach to innovation works. In contrast, services launched with many features or ambitious goals seem always to fail.

Small, measurable successes should be highly valued in the field of persuasive technology. They teach us how to succeed. We can also learn from failure, but those failures should be small and fast. If a big project fails, the design team has taken the wrong path to innovation.

I'm convinced that we can design persuasive technology to do big things. Yes, we can persuade people to stop smoking. We can motivate people to save the environment. I believe we can even leverage persuasive technology to create more harmony among disparate cultures and countries. But we can tackle those big issues effectively only when we—as a community of academics and practitioners—understand how to do the small things well. Many small successes create a foundation of talent and insight for achieving more ambitious goals.