ABOUT THE AUTHORS

Richard H. Thaler is the Ralph and Dorothy Keller Distinguished Service Professor of Behavioral Science and Economics and the director of the Center for Decision Research at the University of Chicago's Graduate School of Business.

Cass R. Sunstein is Felix Frankfurter Professor of Law at Harvard Law School.

NUDGE

Improving Decisions About Health, Wealth, and Happiness

Richard H. Thaler
Cass R. Sunstein

Penguin Books

Edna Ullmann-Margalit, Adrian Vermeule, Eric Wanner, Elke Weber, Roman Weil, Susan Woodward, and Marion Wrobel. As always, our toughest and wisest advice came from France Leclerc and Martha Nussbaum. Vicki Drozd helped out with everything, as she always does, and made sure that all the research assistants got paid, which they appreciated. Thanks too to Ellyn Ruddick-Sunstein, for helpful discussion, patience, both sense and amusement about behavioral economics, and good cheer.

We also owe a special thanks to all the staff at Noodles restaurant on 57th Street. They have fed us and listened to us planning and discussing this book, among other things, for several years now. We'll be back next week.

INTRODUCTION

The Cafeteria

A friend of yours, Carolyn, is the director of food services for a large city school system. She is in charge of hundreds of schools, and hundreds of thousands of kids eat in her cafeterias every day. Carolyn has formal training in nutrition (a master's degree from the state university), and she is a creative type who likes to think about things in nontraditional ways.

One evening, over a good bottle of wine, she and her friend Adam, a statistically oriented management consultant who has worked with supermarket chains, hatched an interesting idea. Without changing any menus, they would run some experiments in her schools to determine whether the way the food is displayed and arranged might influence the choices kids make. Carolyn gave the directors of dozens of school cafeterias specific instructions on how to display the food choices. In some schools the desserts were placed first, in others last, in still others in a separate line. The location of various food items was varied from one school to another. In some schools the French fries, but in others the carrot sticks, were at eye level.

From his experience in designing supermarket floor plans, Adam suspected that the results would be dramatic. He was right. Simply by rearranging the cafeteria, Carolyn was able to increase or decrease the consumption of many food items by as much as 25 percent. Carolyn learned a big lesson: school children, like adults, can be greatly influenced by small

changes in the context. The influence can be exercised for better or for worse. For example, Carolyn knows that she can increase consumption of healthy foods and decrease consumption of unhealthy ones.

With hundreds of schools to work with, and a team of graduate student volunteers recruited to collect and analyze the data, Carolyn believes that she now has considerable power to influence what kids eat. Carolyn is pondering what to do with her newfound power. Here are some suggestions she has received from her usually sincere but occasionally mischievous friends and coworkers:

1. Arrange the food to make the students best off, all things considered.
2. Choose the food order at random.
3. Try to arrange the food to get the kids to pick the same foods they would choose on their own.
4. Maximize the sales of the items from the suppliers that are willing to offer the largest bribes.
5. Maximize profits, period.

Option 1 has obvious appeal, yet it does seem a bit intrusive, even paternalistic. But the alternatives are worse! Option 2, arranging the food at random, could be considered fair-minded and principled, and it is in one sense neutral. But if the orders are randomized across schools, then the children at some schools will have less healthy diets than those at other schools. Is this desirable? Should Carolyn choose that kind of neutrality, if she can easily make most students better off, in part by improving their health?

Option 3 might seem to be an honorable attempt to avoid intrusion: try to mimic what the children would choose for themselves. Maybe that is really the neutral choice, and maybe Carolyn should neutrally follow people's wishes (at least where she is dealing with older students). But a little thought reveals that this is a difficult option to implement. Adam's experiment proves that what kids choose depends on the order in which the items are displayed. What, then, are the true preferences of the children? What does it mean to say that Carolyn should try to figure out what the students would choose "on their own"? In a cafeteria, it is impossible to avoid some way of organizing food.

Option 4 might appeal to a corrupt person in Carolyn's job, and manipulating the order of the food items would put yet another weapon in the arsenal of available methods to exploit power. But Carolyn is honorable and honest, so she does not give this option any thought. Like Options 2 and 3, Option 5 has some appeal, especially if Carolyn thinks that the best cafeteria is the one that makes the most money. But should Carolyn really try to maximize profits if the result is to make children less healthy, especially since she works for the school district?

Carolyn is what we will be calling a *choice architect*. A choice architect has the responsibility for organizing the context in which people make decisions. Although Carolyn is a figment of our imagination, many real people turn out to be choice architects, most without realizing it. If you design the ballot voters use to choose candidates, you are a choice architect. If you are a doctor and must describe the alternative treatments available to a patient, you are a choice architect. If you design the form that new employees fill out to enroll in the company health care plan, you are a choice architect. If you are a parent, describing possible educational options to your son or daughter, you are a choice architect. If you are a salesperson, you are a choice architect (but you already knew that).

There are many parallels between choice architecture and more traditional forms of architecture. A crucial parallel is that there is no such thing as a "neutral" design. Consider the job of designing a new academic building. The architect is given some requirements. There must be room for 120 offices, 8 classrooms, 12 student meeting rooms, and so forth. The building must sit on a specified site. Hundreds of other constraints will be imposed—some legal, some aesthetic, some practical. In the end, the architect must come up with an actual building with doors, stairs, windows, and hallways. As good architects know, seemingly arbitrary decisions, such as where to locate the bathrooms, will have subtle influences on how the people who use the building interact. Every trip to the bathroom creates an opportunity to run into colleagues (for better or for worse). A good building is not merely attractive; it also "works."

As we shall see, small and apparently insignificant details can have major impacts on people's behavior. A good rule of thumb is to assume that "everything matters." In many cases, the power of these small details comes from focusing the attention of users in a particular direction. A wonderful example of this principle comes from, of all places, the men's

rooms at Schiphol Airport in Amsterdam. There the authorities have etched the image of a black housefly into each urinal. It seems that men usually do not pay much attention to where they aim, which can create a bit of a mess, but if they see a target, attention and therefore accuracy are much increased. According to the man who came up with the idea, it works wonders. "It improves the aim," says Aad Kieboom. "If a man sees a fly, he aims at it." Kieboom, an economist, directs Schiphol's building expansion. His staff conducted fly-in-urinal trials and found that etchings reduce spillage by 80 percent.[1]

The insight that "everything matters" can be both paralyzing and empowering. Good architects realize that although they can't build the perfect building, they can make some design choices that will have beneficial effects. Open stairwells, for example, may produce more workplace interaction and more walking, and both of these are probably desirable. And just as a building architect must eventually build some particular building, a choice architect like Carolyn must choose a particular arrangement of the food options at lunch, and by so doing she can influence what people eat. She can nudge.*

Libertarian Paternalism

If, all things considered, you think that Carolyn should take the opportunity to nudge the kids toward food that is better for them, Option

*Please do not confuse *nudge* with *noodge*. As William Safire has explained in his "On Language" column in the *New York Times Magazine* (October 8, 2000), the "Yiddishism *noodge*" is "a noun meaning 'pest, annoying nag, persistent complainer.' . . . To *nudge* is 'to push mildly or poke gently in the ribs, especially with the elbow.' One who *nudges* in that manner—'to alert, remind, or mildly warn another'—is a far *geshrei* from a *noodge* with his incessant, bothersome whining." *Nudge* rhymes with *judge*, while the *oo* sound in *noodge* is pronounced as in *book*.

While we are all down here, a small note about the reading architecture of this book when it comes to footnotes and references. Footnotes such as this one that we deem worth reading are keyed with a symbol and placed at the bottom of the page, so that they are easy to find. We have aimed to keep these to a minimum. Numbered endnotes contain information about source material. These can be skipped by all but the most scholarly of readers. When the authors of cited material are mentioned in the text, we sometimes add a date in parentheses—Smith (1982), for example—to enable readers to go directly to the bibliography without having first to find the endnote.

1, then we welcome you to our new movement: *libertarian paternalism*. We are keenly aware that this term is not one that readers will find immediately endearing. Both words are somewhat off-putting, weighted down by stereotypes from popular culture and politics that make them unappealing to many. Even worse, the concepts seem to be contradictory. Why combine two reviled and contradictory concepts? We argue that if the terms are properly understood, both concepts reflect common sense—and they are far more attractive together than alone. The problem with the terms is that they have been captured by dogmatists.

The libertarian aspect of our strategies lies in the straightforward insistence that, in general, people should be free to do what they like—and to opt out of undesirable arrangements if they want to do so. To borrow a phrase from the late Milton Friedman, libertarian paternalists urge that people should be "free to choose."[2] We strive to design policies that maintain or increase freedom of choice. When we use the term *libertarian* to modify the word *paternalism*, we simply mean liberty-preserving. And when we say liberty-preserving, we really mean it. Libertarian paternalists want to make it easy for people to go their own way; they do not want to burden those who want to exercise their freedom.

The paternalistic aspect lies in the claim that it is legitimate for choice architects to try to influence people's behavior in order to make their lives longer, healthier, and better. In other words, we argue for self-conscious efforts, by institutions in the private sector and also by government, to steer people's choices in directions that will improve their lives. In our understanding, a policy is "paternalistic" if it tries to influence choices in a way that will make choosers better off, *as judged by themselves*.[3] Drawing on some well-established findings in social science, we show that in many cases, individuals make pretty bad decisions—decisions they would not have made if they had paid full attention and possessed complete information, unlimited cognitive abilities, and complete self-control.

Libertarian paternalism is a relatively weak, soft, and nonintrusive type of paternalism because choices are not blocked, fenced off, or significantly burdened. If people want to smoke cigarettes, to eat a lot of candy, to choose an unsuitable health care plan, or to fail to save for retirement, libertarian paternalists will not force them to do otherwise—or even make things hard for them. Still, the approach we recommend does count as pa-

ternalistic, because private and public choice architects are not merely trying to track or to implement people's anticipated choices. Rather, they are self-consciously attempting to move people in directions that will make their lives better. They nudge.

A nudge, as we will use the term, is any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates. Putting the fruit at eye level counts as a nudge. Banning junk food does not.

Many of the policies we recommend can and have been implemented by the private sector (with or without a nudge from the government). Employers, for example, are important choice architects in many of the examples we discuss in this book. In areas involving health care and retirement plans, we think that employers can give employees some helpful nudges. Private companies that want to make money, and to do good, can even benefit from environmental nudges, helping to reduce air pollution (and the emission of greenhouse gases). But as we shall show, the same points that justify libertarian paternalism on the part of private institutions apply to government as well.

Humans and Econs: Why Nudges Can Help

Those who reject paternalism often claim that human beings do a terrific job of making choices, and if not terrific, certainly better than anyone else would do (especially if that someone else works for the government). Whether or not they have ever studied economics, many people seem at least implicitly committed to the idea of *homo economicus*, or economic man—the notion that each of us thinks and chooses unfailingly well, and thus fits within the textbook picture of human beings offered by economists.

If you look at economics textbooks, you will learn that homo economicus can think like Albert Einstein, store as much memory as IBM's Big Blue, and exercise the willpower of Mahatma Gandhi. Really. But the folks that we know are not like that. Real people have trouble with long division if they don't have a calculator, sometimes forget their spouse's birthday,

and have a hangover on New Year's Day. They are not homo economicus; they are homo sapiens. To keep our Latin usage to a minimum we will hereafter refer to these imaginary and real species as Econs and Humans.

Consider the issue of obesity. Rates of obesity in the United States are now approaching 20 percent, and more than 60 percent of Americans are considered either obese or overweight. There is overwhelming evidence that obesity increases risks of heart disease and diabetes, frequently leading to premature death. It would be quite fantastic to suggest that everyone is choosing the right diet, or a diet that is preferable to what might be produced with a few nudges.

Of course, sensible people care about the taste of food, not simply about health, and eating is a source of pleasure in and of itself. We do not claim that everyone who is overweight is necessarily failing to act rationally, but we do reject the claim that all or almost all Americans are choosing their diet optimally. What is true for diets is true for other risk-related behavior, including smoking and drinking, which produce more than five hundred thousand premature deaths each year. With respect to diet, smoking, and drinking, people's current choices cannot reasonably be claimed to be the best means of promoting their well-being. Indeed, many smokers, drinkers, and overeaters are willing to pay third parties to help them make better decisions.

But our basic source of information here is the emerging science of choice, consisting of careful research by social scientists over the past four decades. That research has raised serious questions about the rationality of many judgments and decisions that people make. To qualify as Econs, people are not required to make perfect forecasts (that would require omniscience), but they are required to make unbiased forecasts. That is, the forecasts can be wrong, but they can't be systematically wrong in a predictable direction. Unlike Econs, Humans predictably err. Take, for example, the "planning fallacy"—the systematic tendency toward unrealistic optimism about the time it takes to complete projects. It will come as no surprise to anyone who has ever hired a contractor to learn that everything takes longer than you think, even if you know about the planning fallacy.

Hundreds of studies confirm that human forecasts are flawed and biased. Human decision making is not so great either. Again to take just one example, consider what is called the "status quo bias," a fancy name for in-

ertia. For a host of reasons, which we shall explore, people have a strong tendency to go along with the status quo or default option.

When you get a new cell phone, for example, you have a series of choices to make. The fancier the phone, the more of these choices you face, from the background to the ring sound to the number of times the phone rings before the caller is sent to voice mail. The manufacturer has picked one option as the default for each of these choices. Research shows that whatever the default choices are, many people stick with them, even when the stakes are much higher than choosing the noise your phone makes when it rings.

Two important lessons can be drawn from this research. First, never underestimate the power of inertia. Second, that power can be harnessed. If private companies or public officials think that one policy produces better outcomes, they can greatly influence the outcome by choosing it as the default. As we will show, setting default options, and other similar seemingly trivial menu-changing strategies, can have huge effects on outcomes, from increasing savings to improving health care to providing organs for lifesaving transplant operations.

The effects of well-chosen default options provide just one illustration of the gentle power of nudges. In accordance with our definition, a nudge is any factor that significantly alters the behavior of Humans, even though it would be ignored by Econs. Econs respond primarily to incentives. If the government taxes candy, they will buy less candy, but they are not influenced by such "irrelevant" factors as the order in which options are displayed. Humans respond to incentives too, but they are also influenced by nudges.* By properly deploying both incentives and nudges, we can improve our ability to improve people's lives, and help solve many of society's major problems. And we can do so while still insisting on everyone's freedom to choose.

*Alert readers will notice that incentives can come in different forms. If steps are taken to increase people's cognitive effort—as by placing fruit at eye level and candy in a more obscure place—it might be said that the "cost" of choosing candy is increased. Some of our nudges do, in a sense, impose cognitive (rather than material) costs, and in that sense alter incentives. Nudges count as such, and qualify as libertarian paternalism, only if any costs are low.

A False Assumption and Two Misconceptions

Many people who favor freedom of choice reject any kind of paternalism. They want the government to let citizens choose for themselves. The standard policy advice that stems from this way of thinking is to give people as many choices as possible, and then let them choose the one they like best (with as little government intervention or nudging as possible). The beauty of this way of thinking is that it offers a simple solution to many complex problems: Just Maximize (the number and variety of) Choices—full stop! The policy has been pushed in many domains, from education to prescription drug insurance plans. In some circles, Just Maximize Choices has become a policy mantra. Sometimes the only alternative to this mantra is thought to be a government mandate which is derided as "One Size Fits All." Those who favor Just Maximize Choices don't realize there is plenty of room between their policy and a single mandate. They oppose paternalism, or think they do, and they are skeptical about nudges. We believe that their skepticism is based on a false assumption and two misconceptions.

The false assumption is that almost all people, almost all of the time, make choices that are in their best interest or at the very least are better than the choices that would be made by someone else. We claim that this assumption is false—indeed, obviously false. In fact, we do not think that anyone believes it on reflection.

Suppose that a chess novice were to play against an experienced player. Predictably, the novice would lose precisely because he made inferior choices—choices that could easily be improved by some helpful hints. In many areas, ordinary consumers are novices, interacting in a world inhabited by experienced professionals trying to sell them things. More generally, how well people choose is an empirical question, one whose answer is likely to vary across domains. It seems reasonable to say that people make good choices in contexts in which they have experience, good information, and prompt feedback—say, choosing among ice cream flavors. People know whether they like chocolate, vanilla, coffee, licorice, or something else. They do less well in contexts in which they are inexperienced and poorly informed, and in which feedback is slow or infrequent—say, in choosing between fruit and ice cream (where the long-term effects are

slow and feedback is poor) or in choosing among medical treatments or investment options. If you are given fifty prescription drug plans, with multiple and varying features, you might benefit from a little help. So long as people are not choosing perfectly, some changes in the choice architecture could make their lives go better (as judged by their own preferences, not those of some bureaucrat). As we will try to show, it is not only possible to design choice architecture to make people better off; in many cases it is easy to do so.

The first misconception is that it is possible to avoid influencing people's choices. In many situations, some organization or agent *must* make a choice that will affect the behavior of some other people. There is, in those situations, no way of avoiding nudging in some direction, and whether intended or not, these nudges will affect what people choose. As illustrated by the example of Carolyn's cafeterias, people's choices are pervasively influenced by the design elements selected by choice architects. It is true, of course, that some nudges are unintentional; employers may decide (say) whether to pay employees monthly or biweekly without intending to create any kind of nudge, but they might be surprised to discover that people save more if they get paid biweekly because twice a year they get three pay checks in one month. It is also true that private and public institutions can strive for one or another kind of neutrality—as, for example, by choosing randomly, or by trying to figure out what most people want. But unintentional nudges can have major effects, and in some contexts, these forms of neutrality are unattractive; we shall encounter many examples.

Some people will happily accept this point for private institutions but strenuously object to government efforts to influence choice with the goal of improving people's lives. They worry that governments cannot be trusted to be competent or benign. They fear that elected officials and bureaucrats will place their own interests first, or pay attention to the narrow goals of self-interested private groups. We share these concerns. In particular, we emphatically agree that for government, the risks of mistake, bias, and overreaching are real and sometimes serious. We favor nudges over commands, requirements, and prohibitions in part for that reason. But governments, no less than cafeterias (which governments frequently run), have to provide starting points of one or another kind. This is not avoidable. As we shall emphasize, they do so every day through the rules they

set, in ways that inevitably affect some choices and outcomes. In this respect, the antinudge position is unhelpful—a literal nonstarter.

The second misconception is that paternalism always involves coercion. In the cafeteria example, the choice of the order in which to present food items does not force a particular diet on anyone, yet Carolyn, and others in her position, might select some arrangement of food on grounds that are paternalistic in the sense that we use the term. Would anyone object to putting the fruit and salad before the desserts at an elementary school cafeteria if the result were to induce kids to eat more apples and fewer Twinkies? Is this question fundamentally different if the customers are teenagers, or even adults? Since no coercion is involved, we think that some types of paternalism should be acceptable even to those who most embrace freedom of choice.

In domains as varied as savings, organ donations, marriage, and health care, we will offer specific suggestions in keeping with our general approach. And by insisting that choices remain unrestricted, we think that the risks of inept or even corrupt designs are reduced. Freedom to choose is the best safeguard against bad choice architecture.

Choice Architecture in Action

Choice architects can make major improvements to the lives of others by designing user-friendly environments. Many of the most successful companies have helped people, or succeeded in the marketplace, for exactly that reason. Sometimes the choice architecture is highly visible, and consumers and employers are much pleased by it. (The iPod and the iPhone are good examples because not only are they elegantly styled, but it is also easy for the user to get the devices to do what they want.) Sometimes the architecture is taken for granted and could benefit from some careful attention.

Consider an illustration from our own employer, the University of Chicago. The university, like many large employers, has an "open enrollment" period every November, when employees are allowed to revise the selections they have made about such benefits as health insurance and retirement savings. Employees are required to make their choices online. (Public computers are available for those who would otherwise not have

Internet access.) Employees receive, by mail, a package of materials explaining the choices they have and instructions on how to log on to make these choices. Employees also receive both paper and email reminders.

Because employees are human, some neglect to log on, so it is crucial to decide what the default options are for these busy and absent-minded employees. To simplify, suppose there are two alternatives to consider: those who make no active choice can be given the same choice they made the previous year, or their choice can be set back to "zero." Suppose that last year an employee, Janet, contributed one thousand dollars to her retirement plan. If Janet makes no active choice for the new year, one alternative would be to default her to a one thousand–dollar contribution; another would be to default her to zero contribution. Call these the "status quo" and "back to zero" options. How should the choice architect choose between these defaults?

Libertarian paternalists would like to set the default by asking what reflective employees in Janet's position would actually want. Although this principle may not always lead to a clear choice, it is certainly better than choosing the default at random, or making either "status quo" or "back to zero" the default for everything. For example, it is a good guess that most employees would not want to cancel their heavily subsidized health insurance. So for health insurance the status quo default (same plan as last year) seems strongly preferred to the back to zero default (which would mean going without health insurance).

Compare this to the employee's "flexible spending account," in which an employee sets aside money each month that can be used to pay for certain expenditures (such as uninsured medical or child care expenses). Money put into this account has to be spent each year or it is lost, and the predicted expenditures might vary greatly from one year to the next (for example, child care expenses go down when a child enters school). In this case, the zero default probably makes more sense than the status quo.

This problem is not merely hypothetical. We once had a meeting with three of the top administrative officers of the university to discuss similar issues, and the meeting happened to take place on the final day of the employees' open enrollment period. We mentioned this and asked whether the administrators had remembered to meet the deadline. One said that he was planning on doing it later that day and was glad for the reminder. Another admitted to having forgotten, and the third said that he was hoping that his wife had remembered to do it! The group then turned to the question of what the default should be for a supplementary salary reduction program (a tax-sheltered savings program). To that point, the default had been the "back to zero" option. But since contributions to this program could be stopped at any time, the group unanimously agreed that it would be better to switch to the status quo "same as last year" default. We are confident that many absent-minded professors will have more comfortable retirements as a result.

This example illustrates some basic principles of good choice architecture. Choosers are human, so designers should make life as easy as possible. Send reminders, and then try to minimize the costs imposed on those who, despite your (and their) best efforts, space out. As we will see, these principles (and many more) can be applied in both the private and public sectors, and there is much room for going beyond what is now being done.

A New Path

We shall have a great deal to say about private nudges. But many of the most important applications of libertarian paternalism are for government, and we will offer a number of recommendations for public policy and law. Our hope is that that those recommendations might appeal to both sides of the political divide. Indeed, we believe that the policies suggested by libertarian paternalism can be embraced by Republicans and Democrats alike. A central reason is that many of those policies cost little or nothing; they impose no burden on taxpayers at all.

Many Republicans are now seeking to go beyond simple opposition to government action. As the experience with Hurricane Katrina showed, government is often required to act, for it is the only means by which the necessary resources can be mustered, organized, and deployed. Republicans want to make people's lives better; they are simply skeptical, and legitimately so, about eliminating people's options.

For their part, many Democrats are willing to abandon their enthusiasm for aggressive government planning. Sensible Democrats certainly hope that public institutions can improve people's lives. But in many domains, Democrats have come to agree that freedom of choice is a good and even

indispensable foundation for public policy. There is a real basis here for crossing partisan divides.

Libertarian paternalism, we think, is a promising foundation for bipartisanship. In many domains, including environmental protection, family law, and school choice, we will be arguing that better governance requires less in the way of government coercion and constraint, and more in the way of freedom to choose. If incentives and nudges replace requirements and bans, government will be both smaller and more modest. So, to be clear: *we are not for bigger government, just for better governance.*

Actually we have evidence that our optimism (which we admit may be a bias) is more than just rosy thinking. Libertarian paternalism with respect to savings, discussed in Chapter 6, has received enthusiastic and widespread bipartisan support in Congress, including from current and former conservative Republican senators such as Robert Bennett (Utah) and Rick Santorum (Pa.) and liberal Democrats such as Rahm Emanuel of Illinois. In 2006 some of the key ideas were quietly enacted into law. The new law will help many Americans have more comfortable retirements but costs essentially nothing in taxpayer dollars.

In short, libertarian paternalism is neither left nor right, neither Democratic nor Republican. In many areas, the most thoughtful Democrats are going beyond their enthusiasm for choice-eliminating programs. In many areas, the most thoughtful Republicans are abandoning their knee-jerk opposition to constructive governmental initiatives. For all their differences, we hope that both sides might be willing to converge in support of some gentle nudges.

PART I

HUMANS AND ECONS

13

IMPROVING SCHOOL CHOICES

In 1944 President Franklin Delano Roosevelt included "the right to a good education" in what he called a Second Bill of Rights, designed to promote "security" and suitable for a modern democracy.[1] Most Americans seem to believe that children do have a right to a good education; there is a consensus on that point. One reason for that consensus is that educated people are more free. But the consensus breaks down when people explore how, exactly, to achieve that right.

School choice remains an intensely polarizing issue in American politics. The case for choice was originally popularized by the great libertarian economist Milton Friedman. His argument is a simple one: the best way to improve our children's schools is to introduce competition. If schools compete, kids win. And if schools compete, those who are the least advantaged have the most to gain. Wealthy families already have "school choice," because they can send their children to private schools. If we give parents vouchers to send their children to any school they want, then we will put children from poor families more nearly on a par with their more privileged middle- and upper-class counterparts. Shouldn't poor children have the same rights that wealthy ones do?

Critics of school choice argue that such programs amount, in practice, to an attack on the public school system that has helped make America great. The critics worry that in the end, public schools, which serve diverse people and allow them to be educated together, will lose both students and money. They fear that vouchers will turn out to be a subsidy to rich

parents who can already afford to send their children to fancy private schools—and even worse, that public schools will end up with the kids that the private schools don't want.

As libertarians, we are strongly inclined to support the concept of school choice, because freedom is usually a good idea and because competition is likely to improve education. But an abstract preference for choice does not allow us to select any particular plan, and of course the proof is in the pudding. We have seen that the Just Maximize Choices mantra does not always lead to the best possible outcomes. So we need to ask, when it comes to schools, do more choices actually help? Since the 1970s cities around the country have experimented with choice programs, providing observers with the chance to assess the actual effects of such programs. The evidence suggests that while choice programs are hardly a panacea, they can indeed improve student performance. Carolyn Hoxby, a leading economist who has analyzed both voucher and charter school programs, finds that when facing competition, public schools produce higher student achievement per dollar spent. Test-score improvements can range from 1 to 7 percent a year depending on the school and student—and improvement is usually greatest among younger students, low-income students, and minority-group members.[2]

Even though the results suggest that school choice can and does help, we believe that the results could be significantly enhanced by helping parents make better choices on behalf of their children. Many parents simply do not make use of their options and instead just send their child to the default school (usually, but not always, their neighborhood school). And those who do make choices are sometimes ill prepared to make good ones. Because we approve of more choice, we want to focus on one important part of the school choice issue—how to create plans that put parents in a position to make sensible decisions for their children.

Complex Choices and Mental Shortcuts

Consider the revealing case of Worcester, Massachusetts. President Bush signed the federal No Child Left Behind law in 2001, with the goal of increasing public school accountability by mandating certain testing standards. (We put to one side the many controversial questions raised by that law.) By June 2003 twelve of Worcester's fifty public schools had been labeled "in need of improvement" for two consecutive years, and five for three consecutive years. That summer, forty-seven hundred students, almost one-fifth the district's student population, were eligible to transfer, and eighteen hundred students had the right to collect federal money for supplemental education services. But six months later, only one student had switched schools, and only two had taken advantage of supplemental services!

Worcester officials themselves were primarily responsible. True, the school system notified parents at underperforming schools about their rights under No Child Left Behind. But it also engaged in what the political scientist William Howell calls "friendly discouragement," making parents reluctant to exercise their right to choose.[3] The school system qualified its language about the meaning of *underperforming*, stressed the limitations of the No Child Left Behind evaluation criteria, and highlighted unattractive parts of No Child Left Behind, noting that space limitations might not permit transfers to be processed. The school system also explained that it was trying to improve.

For the undeterred, exercising choice was a tedious, multistage process. First, parents had to meet with their school's principal. Few did. Next, they had to attend another meeting at a school information center. The center's director said that two parents expressed interest in such a meeting. At these meetings, district officials again stressed that transfers were not always possible and that there were no guarantees about transportation or school location. And all of that was before parents had to file the transfer paperwork. Even worse, because the school district controlled access to information, tutoring service and test prep companies could not reach students without the district's blessing. The companies essentially depended on positive comments from the school district.

As with a 401(k) plan, the average parents know little about their child's school, let alone all the other schools that are available. They might well stick with the status quo or ultimately make poor decisions. The trick is to promote actual freedom—not just by giving people lots of choices (though that can help) but also by putting people in a good position to choose what would be best for their children. Consider a few details.

When parents pick schools, status quo bias plays a big role. The neigh-

borhood school that one knows, failing or not, may be preferable to the unknown school half an hour away. In any case, the Byzantine nature of collecting and distributing school data makes it difficult for parents to think through their options. In Charlotte, North Carolina, for instance, parents receive a hundred-page booklet with descriptions of 190 schools written by representatives of the schools themselves, emphasizing each school's positive features. The booklet does *not* include information on physical locations, test scores, attendance rates, and racial composition—these are available only on the district Web site. Meanwhile, staff members at a special district-wide application center are instructed to respond to questions like "Which school is the best school?" by saying that "a good school depends on each individual child" and advising parents to talk to their children about what their needs are, and to visit the various schools in order to determine which is best for their children. Although this advice is unobjectionable, it is about as helpful as when a waiter responds to an inquiry about what is good by saying: "Everything!"

A creative experiment in Charlotte shows that choices can be improved with better and simpler information.[4] Charlotte gave parents the option to apply for admission at multiple public schools besides their default school. Low-income parents tended to put less weight than high-income parents on school quality, as measured by test scores, and rarely tried to enroll in higher-performing schools. A random sample of parents was selected to receive an abbreviated "fact sheet" about the schools—much in the spirit of the RECAP idea that we have suggested in other areas. Printed on each sheet was a complete listing of average test scores and acceptance rates, from highest to lowest, at schools available to a given child.

The experimenters wanted to find out whether parents, and especially low-income parents, would choose better schools. They did. Much better ones. The parents who received the fact sheets made decisions implying that the weight they assigned to school quality (as measured by test scores) had doubled. The schools they selected had, on average, 70 percent higher test scores than the scores at their neighborhood schools. This had the effect of making their choices similar to those of families whose incomes were $65,000 a year higher. Furthermore, when children are lucky enough to switch to better schools, their performance improves considerably. The students who are lucky enough to win the lotteries held to decide who gets to attend the popular better schools are less likely to be suspended and have higher test scores than the students who lost.[5]

Incentive Conflicts and Matching

A good choice architect can do more than help parents achieve what is already in their own self-interest. The architect can also help reduce latent incentive conflicts between advantaged and disadvantaged parents during the choice process.

Despite the attention they receive in the media, market-based programs like vouchers are available to relatively few students nationwide. One popular alternative is a policy known as controlled choice, which emerged in the wake of 1970s court rulings prohibiting busing for the purpose of achieving desegregation. The idea was to continue integration by guaranteeing students a priority space at a nearby school or a school that a sibling attended, while giving them the option to apply for enrollment somewhere else.

School administrators in Boston adopted a computer algorithm designed to assign as many students as possible to their first-choice schools, while still giving priority to the neighborhood students. It is hard to know exactly how many districts use the so-called Boston system, because administrators do not always explain controlled-choice policies in detail, but some of the larger metropolitan districts that employ that algorithm or something similar include Denver, Tampa, Minneapolis, Louisville, and Seattle. (If two students applied to a school with one open seat, Seattle and Louisville broke the tie on the basis of race, a practice the Supreme Court ruled unconstitutional in 2007.)

Matching as many first choices as possible sounds sensible enough, except for one problem. Picking schools in the Boston system turns out to be a complex game of strategy, with the winners reaping the spoils. How do the winners win? They lie, a little. Economists call it strategic misrepresentation.

There is a mathematical (and complicated) reason why lying is a good strategy in the Boston system, but to get an intuitive feel imagine that college admissions suddenly operated on a national controlled-choice system. Schools like Harvard and Stanford would be heavily overdemanded, and

locals would get preferential treatment. You would have only slightly better odds of getting into one than of winning the Powerball jackpot. (You think property in Cambridge and Palo Alto is expensive now? What if living there guaranteed your child a seat at Harvard or Stanford?) Clever parents who do not happen to live in Cambridge, but who have been dreaming of sending their child to Harvard since the diaper days, would realize the futility of listing it first. The Boston system attempts to match as many first choices as possible, so if every honest parent in America listed Harvard first, only Cambridge residents could sleep well at night.

Instead of taking their chances on a long shot, parents outside Cambridge would be better served to select as their first choice a slightly less popular school such as Dartmouth or Cornell, say, where there are also fewer students nearby getting preferential treatment. In the Boston system, parents who rank a school second or third lose out to everyone who ranks it first—making it risky to use a first choice on a highly sought-after school if a child has a low priority, and a complete waste to list such a school as a second choice. Information about school demand is usually available online, giving parents an incentive to tweak rankings based on acceptance rates and where their child has priority.

When the Boston system was first developed, almost no one intuited this strategy. (Only a handful of people even knew how the algorithm worked!) But over time, some parents figured out ways to gain an edge. Not surprisingly, affluent, educated parents with large social networks (they volunteer at school with other affluent, educated parents) learned the tricks first. They performed better than less affluent, less educated parents, who routinely listed an overdemanded school as a second choice, the worst mistake they could make. Who knows how many of their children lost out on access to first-rate educations because of it?

The Boston system is still in place around the country, though not in Boston. In 2003 a group of economists led by Al Roth at Harvard pointed out these problems to initially skeptical Boston school administrators. After letting the economists poke around in the internal data, the administrators became convinced of their system's flaws.[6]

In response, they adopted the economists' new strategy-proof choice mechanism, based on one used to match hospitals and medical residents.

The mechanism does not penalize parents who are unsophisticated about the choice process, allowing them to spend time visiting schools and seeing teachers, rather than estimating the level of competition to get into each school. In return, administrators do not have to guess about parents' true preferences so that the policy can be adjusted properly based on future feedback.

Nudging High Schoolers Toward College

Good choice architecture doesn't need to originate with a wonkish professor and a powerful computer algorithm. It can be the brainchild of a local school official or two. In San Marcos, Texas, the school superintendent and an administrator at nearby Austin Community College were looking for a way to get more of San Marcos's largely Latino student population into college. They hit on a nudge so simple and effective it spread through the state faster than a YouTube clip. (Well, maybe not that fast.) The nudge was this: in order to graduate from San Marcos High, a student would have to complete an application to nearby Austin Community College. Because all it takes to get admitted to the community college is a high school degree and a record of having taken a standardized test, completing the application properly was tantamount to acceptance.

In San Marcos, schools run on a tight budget, and two-thirds of high schoolers never experience higher education. The superintendent had no outside funding to implement the idea, so she asked her teachers and the community college for help. Students were pulled from English classes to meet with the college's staff counselors. In a smart piece of mapping, the counselors didn't try to sell the students on the high-mindedness of education. Instead, they hooked them with the universal symbol of teenage freedom: the automobile. They talked about how much more money college graduates earned compared with high school graduates, explaining it as the difference between a Mercedes and a KIA. Next, community college administrators took a standardized admissions exam to the high school and tested the students free of charge. The administrators also gave students financial aid information and had tax consultants offer weekend sessions for parents.

In the end, the nudge produced big results. From 2004 to 2005 the percentage of San Marcos High students who went to Texas colleges rose 11 percentage points, to 45 percent. Now more than forty-five Texas high schools have similar programs, and schools in Florida and California have created programs modeled after San Marcos's. In Maine a state legislator is proposing a law requiring high school seniors to submit at least one college application before they graduate.

We have covered a lot of territory in a short space. Milton Friedman was right: at least in the abstract, school choice is an excellent idea, because it increases freedom and offers real promise for improving education. Of course, reforms should be assessed empirically, not in the abstract. Fortunately, existing evidence suggests that school choice has considerable promise.

The major problem, and our principal concern here, is that what is true for investments and prescription drugs is true for education as well: it is not enough to make lots of choices available and then hope parents choose wisely. School systems need to put parents in a position to think through their choices, and to exercise their freedom rather than to rely on the default option. Both parents and children need the right incentives. FDR's "right to a good education" is not part of the Constitution, but it has become a cultural commitment, and a few simple steps could enable many more children to enjoy that right.

14

SHOULD PATIENTS BE FORCED TO BUY LOTTERY TICKETS?

Every election cycle, presidential contenders unveil plans to make health care coverage available to the tens of millions of Americans who lack health insurance. The candidates decry our government's failure, thus far, to implement an effective plan.

Whatever happens in the long run, such plans are hard to design for a simple reason: health care is really expensive. It is expensive in part because Americans want access to all the best services: doctors, hospitals, prescription drugs, medical devices, and nursing homes, to name a few.

Of course, we can try to keep health care affordable on our own, by maintaining healthy lifestyles, and by buying only the health care products and services that we need. We can save money by visiting the doctor no more often than necessary, and if we purchase insurance, we can choose a plan that covers only catastrophic illnesses instead of coverage with low deductibles, which is much more expensive. But there is something that every health care customer in America is forced to buy, whether she wants it or not: the right to sue the doctor for negligence.

Our principal claim here is that patients and doctors should be free to make their own agreements about that right. If patients want to waive the right to sue, they should be allowed to do exactly that. This increase in freedom is likely to help doctors and patients alike, and to make a valuable, even if modest, contribution to the health care problem.

It may seem strange to think that we "purchase" the right to sue. Of course, that right is not an itemized portion of the insurance bill—but it is