


A N E W Y O R K T I M E S B e s t s e l l e r

A BEST BOOK  
OF THE YEAR  
The  
Economist



# *Nudge*

Improving Decisions About  
Health, Wealth, and Happiness

Richard H. Thaler and Cass R. Sunstein

*Revised and Expanded Edition*

"One of the few books I've read recently that fundamentally changes the way I think about the world." —Steven D. Levitt, coauthor of *Freakonomics*



cate it. When many people were still afraid of flying, it was common to see airline flight insurance sold at airports at exorbitant prices. There were no booths in airports selling people advice not to buy such insurance.

In many markets, firms will be competing for the same consumers but will be offering products that are not merely different but that directly oppose each other. Some firms sell cigarettes; others sell products that help you quit smoking. Some firms sell fast food; others sell diet advice. If all consumers are Econs, then there is no reason to worry about which of these competing interests wins. But if some of the consumers are Humans who sometimes make bad choices (as judged by themselves, of course), then all of us may have an interest in which set of firms wins the battle. Government can, of course, outlaw some kinds of activities, but as libertarian paternalists we prefer to nudge—and we are keenly aware that governments are populated by Humans.

What can be done to help? In the next chapter we describe our primary tool: choice architecture.

# 5

## CHOICE ARCHITECTURE

Early in Thaler's career, he was teaching a class on managerial decision making to business school students. Students would sometimes leave class early to go for job interviews (or a golf game) and would try to sneak out of the room as surreptitiously as possible. Unfortunately for them, the only way out of the room was through a large double door in the front, in full view of the entire class (though not directly in Thaler's line of sight). The doors were equipped with large, handsome wood handles, vertically mounted cylindrical pulls about two feet in length. When the students came to these doors, they were faced with two competing instincts. One instinct says that to leave a room you push the door. The other instinct says, when faced with large wooden handles that are obviously designed to be grabbed, you pull. It turns out that the latter instinct trumps the former, and every student leaving the room began by pulling on the handle. Alas, the door opened outward.

At one point in the semester, Thaler pointed this out to the class, as one embarrassed student was pulling on the door handle while trying to escape the classroom. Thereafter, as a student got up to leave, the rest of the class would eagerly wait to see whether the student would push or pull. Amazingly, most still pulled! Their Automatic Systems triumphed; the signal emitted by that big wooden handle simply could not be screened out. (And when Thaler would leave that room on other occasions, he sheepishly found himself pulling too.)

Those doors are bad architecture because they violate a simple psycho-

logical principle with a fancy name: stimulus response compatibility. The idea is that you want the signal you receive (the stimulus) to be consistent with the desired action. When there are inconsistencies, performance suffers and people blunder.

Consider, for example, the effect of a large, red, octagonal sign that said GO. The difficulties induced by such incompatibilities are easy to show experimentally. One of the most famous such demonstrations is the Stroop (1935) test. In the modern version of this experiment people see words flashed on a computer screen and they have a very simple task. They press the right button if they see a word that is displayed in red, and press the left button if they see a word displayed in green. People find the task easy and can learn to do it very quickly with great accuracy. That is, until they are thrown a curve ball, in the form of the word GREEN displayed in red, or the word RED displayed in green. For these incompatible signals, response time slows and error rates increase. A key reason is that the Automatic System reads the word faster than the color naming system can decide the color of the text. See the word GREEN in red text and the nonthinking Automatic System rushes to press the left button, which is, of course, the wrong one. You can try this for yourself. Just get a bunch of colored crayons and write a list of color names, making sure that most of the names are not the same as the color they are written in. (Better yet, get a nearby kid to do this for you.) Then name the color names as fast as you can (that is, read the words and ignore the color): easy, isn't it? Now say the color that the words are written in as fast as you can and ignore the word itself: hard, isn't it? In tasks like this, Automatic Systems always win over Reflective ones.

Although we have never seen a green stop sign, doors such as the ones described above are commonplace, and they violate the same principle. Flat plates say "push me" and big handles say "pull me," so don't expect people to push big handles! This is a failure of architecture to accommodate basic principles of human psychology. Life is full of products that suffer from such defects. Isn't it obvious that the largest buttons on a television remote control should be the power, channel, and volume controls? Yet how many remotes do we see that have the volume control the same size as the "input" control button (which if pressed accidentally can cause the picture to disappear)?

It is possible, however, to incorporate human factors into design, as


Don Norman's wonderful book *The Design of Everyday Things* (1990) illustrates. One of his best examples is the design of a basic four-burner stove (Figure 5.1). Most such stoves have the burners in a symmetric arrangement, as in the stove pictured at the top, with the controls arranged in a linear fashion below. In this set-up, it is easy to get confused about which knob controls the front burner and which controls the back, and many pots and pans have been burned as a result. The other two designs we have illustrated are only two of many better possibilities.

Norman's basic lesson is that designers need to keep in mind that the users of their objects are Humans who are confronted every day with myriad choices and cues. The goal of this chapter is to develop the same idea for choice architects. If you indirectly influence the choices other people make, you are a choice architect. And since the choices you are influencing are going to be made by Humans, you will want your architecture to reflect a good understanding of how humans behave. In particular, you will want to ensure that the Automatic System doesn't get all confused. In this chapter, we offer some basic principles of good (and bad) choice architecture.

### Defaults: Padding the Path of Least Resistance

For reasons we have discussed, many people will take whatever option requires the least effort, or the path of least resistance. Recall the discussion of inertia, status quo bias, and the "yeah, whatever" heuristic. All these forces imply that if, for a given choice, there is a default option—an option that will obtain if the chooser does nothing—then we can expect a large number of people to end up with that option, whether or not it is good for them. And as we have also stressed, these behavioral tendencies toward doing nothing will be reinforced if the default option comes with some implicit or explicit suggestion that it represents the normal or even the recommended course of action.

Defaults are ubiquitous and powerful. They are also unavoidable in the sense that for any node of a choice architecture system, there must be an associated rule that determines what happens to the decision maker if she does nothing. Of course, usually the answer is that if I do nothing, nothing changes; whatever is happening continues to happen. But not always.



5.1  
Three designs of four-burner stovetops

Some dangerous machines, such as chain saws and lawn mowers, are designed with “dead man switches,” so that once you are no longer gripping the machine, it stops. When you leave your computer alone for a while to answer a phone call, nothing is likely to happen until you have talked for a long time, at which point the screen saver comes on, and if you neglect the computer long enough, it may lock itself.

Of course, you can choose how long it takes before your screen saver comes on, but implementing that choice takes some action. Your computer probably came with a default time lag and a default screen saver. Chances are, those are the settings you still have.

Many organizations in both the public and the private sector have discovered the immense power of default options. Successful businesses certainly have. Remember the idea of automatic renewal for magazine subscriptions? If renewal is automatic, many people will subscribe, for a long time, to magazines they don’t read. Business offices at most magazines are aware of that fact. When you download a new piece of software, you will often have numerous choices to make. Do you want the “regular” or “custom” installation? Normally, one of the boxes is already checked, indicating it is the default. Which boxes do the software suppliers check? Two different motives are readily apparent: helpful and self-serving. In the helpful category would be making the regular installation the default if most users will have trouble with the custom installation. In the self-serving category would be making the default a willingness to receive emails with information about new products. In our experience, most software comes with helpful defaults regarding the type of installation, but many come with self-serving defaults on other choices. We will have more to say about motives later. For now, note that not all defaults are selected to make the chooser’s life easier or better.

The choice of the default can be quite controversial. Here is one example. An obscure portion of the No Child Left Behind Act requires that school districts supply the names, addresses, and telephone numbers of students to the recruiting offices of branches of the armed forces. However, the law stipulates that “a secondary school student or the parent of the student may request that the student’s name, address, and telephone listing not be released without prior written parental consent, and the local educational agency or private school shall notify parents of the option

to make a request and shall comply with any request.” Some school districts, such as Fairport, New York, interpreted this law as allowing them to implement an “opt-in” policy. That is, parents were notified that they could elect to make their children’s contact information available, but if they did not do anything, this information would be withheld.

This reading of the law did not meet with the approval of then-Secretary of Defense Donald Rumsfeld. The Defense and Education Departments sent a letter to school districts asserting that the law required an opt-out implementation. Only if parents actively requested that the contact information on their children be withheld would that option apply. In typical bureaucratic language, the departments contended that the relevant laws “do not permit LEA’s [local educational agencies] to institute a policy of *not* providing the required information unless a parent has affirmatively agreed to provide the information.”<sup>1</sup> Both the Defense Department and the school districts realized that opt-in and opt-out policies would lead to very different outcomes. Not surprisingly, much hue and cry ensued. We discuss a similarly touchy subject involving defaults in our chapter on organ donations.

We have emphasized that default rules are inevitable—that private institutions and the legal system cannot avoid choosing them. In some cases, though not all, there is an important qualification to this claim. The choice architect can force the choosers to make their own choice. We call this approach “required choice” or “mandated choice.” In the software example, required choice would be implemented by leaving all the boxes unchecked, and by requiring that at every opportunity one of the boxes be checked in order for people to proceed. In the case of the provision of contact information to the military recruiters, one could imagine a system in which all students (or their parents) are required to fill out a form indicating whether they want to make their contact information available. For emotionally charged issues like this one, such a policy has considerable appeal, because people might not want to be defaulted into an option that they might hate (but fail to reject because of inertia, or real or apparent social pressure).

We believe that required choice, favored by many who like freedom, is sometimes the best way to go. But consider two points about that approach. First, Humans will often consider required choice to be a nuisance

or worse, and would much prefer to have a good default. In the software example, it is really helpful to know what the recommended settings are. Most users do not want to have to read an incomprehensible manual in order to determine which arcane setting to elect. When choice is complicated and difficult, people might greatly appreciate a sensible default. It is hardly clear that they should be forced to choose.

Second, required choosing is generally more appropriate for simple yes-or-no decisions than for more complex choices. At a restaurant, the default option is to take the dish as the chef usually prepares it, with the option to ask that certain ingredients be added or removed. In the extreme, required choosing would imply that the diner has to give the chef the recipe for every dish she orders! When choices are highly complex, required choosing may not be a good idea; it might not even be feasible.

### Expect Error

Humans make mistakes. A well-designed system expects its users to err and is as forgiving as possible. Some examples from the world of real design illustrate this point:

- In the Paris subway system, Le Métro, users insert a paper card the size of a movie ticket into a machine that reads the card, leaves a record on the card that renders it “used,” and then spits it out from the top of the machine. The cards have a magnetic strip on one side but are otherwise symmetric. On Thaler’s first visit to Paris, he was not sure how to use the system, so he tried putting the card in with the magnetic strip face up and was pleased to discover that it worked. He was careful thereafter to insert the card with the strip face up. Many years and trips to Paris later, he was proudly demonstrating to a visiting friend the correct way to use the Metro system when his wife started laughing. It turns out that it doesn’t matter which way you put the card into the machine!

In stark contrast to Le Métro is the system used in most Chicago parking garages. When entering the garage, you put your credit card into a machine that reads it and remembers you. Then when leaving, you must insert the card again into another machine at the exit. This involves reaching out of the car window and inserting the card into a slot. Because credit cards are not symmetric, there are four possible ways to put the card into

the slot (face up or down, strip on the right or left). Exactly one of those ways is the right way. And in spite of a diagram above the slot, it is very easy to put the card in the wrong way, and when the card is spit back out, it is not immediately obvious what caused the card to be rejected or to recall which way it was inserted the first time. Both of us have been stuck for several painful minutes behind some idiot who was having trouble with this machine, and have to admit to having occasionally been the idiot that is making all the people behind him start honking.

- Over the years, automobiles have become much friendlier to their Human operators. If you do not buckle your seat belt, you are buzzed. If you are about to run out of gas, a warning sign appears and you might be beeped. If you need an oil change, your car might tell you. Many cars come with an automatic switch for the headlights that turns them on when you are operating the car and off when you are not, eliminating the possibility of leaving your lights on overnight and draining the battery.

But some error-forgiving innovations are surprisingly slow to be adopted. Take the case of the gas tank cap. On any sensible car the gas cap is attached by a piece of plastic, so that when you remove the cap you cannot possibly drive off without it. Our guess is that this bit of plastic cannot cost more than ten cents. Once some firm had the good idea to include this feature, what excuse can there ever have been for building a car without one?

Leaving the gas cap behind is a special kind of predictable error psychologists call a "postcompletion" error.<sup>2</sup> The idea is that when you have finished your main task, you tend to forget things relating to previous steps. Other examples include leaving your ATM card in the machine after getting your cash, or leaving the original in the copying machine after getting your copies. Most ATMs (but not all) no longer allow this error because you get your card back immediately. Another strategy, suggested by Norman, is to use what he calls a "forcing function," meaning that in order to get what you want, you have to do something else first. So if in order to get your cash, you have to remove the card, you will not forget to do so.

- Another automobile-related bit of good design involves the nozzles for different varieties of gasoline. The nozzles that deliver diesel fuel are too large to fit into the opening on cars that use gasoline, so it is not possi-

ble to make the mistake of putting diesel fuel in your gasoline-powered car (though it is still possible to make the opposite mistake). The same principle has been used to reduce the number of errors involving anesthesia. One study found that human error (rather than equipment failure) caused 82 percent of the "critical incidents." A common error was that the hose for one drug was hooked up to the wrong delivery port, so the patient received the wrong drug. This problem was solved by designing the equipment so that the gas nozzles and connectors were different for each drug. It became physically impossible to make this previously frequent mistake.<sup>3</sup>

- A major problem in health care is called "drug compliance." Many patients, especially the elderly, are on medicines they must take regularly, and in the correct dosage. So here is a choice-architecture question. If you are designing a drug, and you have complete flexibility, how often would you want your patients to have to take their medicine?

If we rule out a one-time dose administered immediately by the doctor (which would be best on all dimensions but is often technically infeasible), then the next-best solution is a medicine taken once a day, preferably in the morning. It is clear why once a day is better than twice (or more) a day, because the more often you have to take the drug, the more opportunities you have to forget. But frequency is not the only concern; regularity is also important. Once a day is much better than once every other day, because the Automatic System can be educated to think: "My pill(s) every morning, when I wake up." Taking the pill becomes a habit, and habits are controlled by the Automatic System. By contrast, remembering to take your medicine every other day is beyond most of us. (Similarly, meetings that occur every week are easier to remember than those that occur every other week.) Some medicines are taken once a week, and most patients take this medicine on Sundays (because that day is different from other days for most people and thus easy to associate with taking one's medicine).

Birth control pills present a special problem along these lines, because they are taken every day for three weeks and then skipped for one week. To solve this problem and to make the process automatic, the pills are typically sold in a special container that contains twenty-eight pills, each in a numbered compartment. Patients are instructed to take a pill every day, in order. The pills for days twenty-two through twenty-eight are placebos whose only role is to facilitate compliance for Human users.

- While working on this book, Thaler sent an email to his economist friend Hal Varian, who is affiliated with Google. Thaler intended to attach a draft of the introduction to give Hal a sense of what the book was about, but forgot the attachment. When Hal wrote back to ask for the missing attachment, he noted with pride that Google was experimenting with a new feature on its email program “gmail” that would solve this problem. A user who mentions the word *attachment* but does not include one would be prompted, “Did you forget your attachment?” Thaler sent the attachment along and told Hal that this was exactly what the book was about.

- Visitors to London who come from the United States or Europe have a problem being safe pedestrians. They have spent their entire lives expecting cars to come at them from the left, and their Automatic System knows to look that way. But in the United Kingdom automobiles drive on the left-hand side of the road, and so the danger often comes from the right. Many pedestrian accidents occur as a result. The city of London tries to help with good design. On many corners, especially in neighborhoods frequented by tourists, the pavement has signs that say, “Look right!”

### Give Feedback

The best way to help Humans improve their performance is to provide feedback. Well-designed systems tell people when they are doing well and when they are making mistakes. Some examples:

- Digital cameras generally provide better feedback to their users than film cameras. After each shot, the photographer can see a (small) version of the image just captured. This eliminates all kinds of errors that were common in the film era, from failing to load the film properly (or at all), to forgetting to remove the lens cap, to cutting off the head of the central figure of the picture. However, early digital cameras failed on one crucial feedback dimension. When a picture was taken, there was no audible cue to indicate that the image had been captured. Modern models now include a very satisfying but completely fake “shutter click” sound when a picture has been taken. (Some cell phones, aimed at the elderly, include a fake dial tone, for similar reasons.)

- An important type of feedback is a warning that things are going wrong, or, even more helpful, are about to go wrong. Our laptops warn us

to plug in or shut down when the battery is dangerously low. But warning systems have to avoid the problem of offering so many warnings that they are ignored. If our computer constantly nags us about whether we are sure we want to open that attachment, we begin to click “yes” without thinking about it. These warnings are thus rendered useless.

- The Department of Homeland Security’s color-coded terror alert system is a nice illustration of feedback that would be useless even if it weren’t incessant. When walking through an American airport any time since 2002, one is bound to hear the following announcement: “The Department of Homeland Security has raised the National Threat Advisory to Orange.” Aside from putting our toiletries into a one-quart zip-lock bag, exactly what actions are we expected to take as a result of this warning?

A look at the Homeland Security Web site provides the answer. We are told: “All Americans should continue to be vigilant, take notice of their surroundings, and report suspicious items or activities to local authorities immediately.” Weren’t we supposed to be doing this at level Yellow? It is a safe bet that these announcements are useless. (Much more useful would be a supply of one-quart zip-lock bags for absentminded travelers; and many airports do in fact provide these.)

- Feedback can be improved in many activities. Consider the simple task of painting a ceiling. This task is more difficult than it might seem because ceilings are nearly always painted white, and it can be hard to see exactly where you have painted. Later, when the paint dries, the patches of old paint will be annoyingly visible. How to solve this problem? Some helpful person invented a type of ceiling paint that goes on pink when wet but turns white when dry. Unless the painter is so colorblind that he can’t tell the difference between pink and white, this solves the problem.

### Understanding “Mappings”: From Choice to Welfare

Some tasks are easy, like choosing a flavor of ice cream; other tasks are hard, like choosing a medical treatment. Consider, for example, an ice cream shop where the varieties differ only in flavor, not calories or other nutritional content. Selecting which ice cream to eat is merely a matter of choosing the one that tastes best. If the flavors are all familiar, such as vanilla, chocolate, and strawberry, most people will be able to predict with

considerable accuracy the relation between their choice and their ultimate consumption experience. Call this relation between choice and welfare a mapping. Even if there are some exotic flavors, the ice cream store can solve the mapping problem by offering a free taste.

Choosing among treatments for some disease is quite another matter. Suppose you are told that you have been diagnosed with prostate cancer and must choose among three options: surgery, radiation, and "watchful waiting" (which means do nothing for now). Each of these options comes with a complex set of possible outcomes regarding side effects of treatment, quality of life, length of life, and so forth. Comparing the options involves making such trade-offs as the following: Would I be willing to risk a one-third chance of impotence or incontinence in order to increase my life expectancy by 3.2 years? This is a hard decision at two levels. First, the patient is unlikely to know these trade-offs, and second, he is unlikely to be able to imagine what life would be like if he were incontinent. Yet here are two scary facts about this scenario. First, most patients decide which course of action to take in the very meeting at which their doctor breaks the bad news about the diagnosis. Second, the treatment option they choose depends strongly on the type of doctor they see.<sup>4</sup> (Some specialize in surgery, others in radiation. None specialize in watchful waiting. Guess which option we suspect might be underutilized?)

The comparison between ice cream and treatment options illustrates the concept of mapping. A good system of choice architecture helps people to improve their ability to map and hence to select options that will make them better off. One way to do this is to make the information about various options more comprehensible, by transforming numerical information into units that translate more readily into actual use. If I am buying apples to make into apple cider, it helps to know the rule of thumb that it takes three apples to make one glass of cider.

Take the example of choosing a digital camera. Cameras advertise their megapixels, and the impression created is certainly that the more megapixels the better. This assumption is itself subject to question, because photos taken with more megapixels take up more room on the camera's storage device and a computer's hard drive. But what is really problematic for consumers is translating megapixels (not the most intuitive concept) into what they care about. Is it worth paying an additional hundred dollars to go

from four to five megapixels? Suppose instead that manufacturers listed the largest print size recommended for a given camera. Instead of being given the options of three, five, or seven megapixels, consumers might be told that the camera can produce quality photos at  $4 \times 6$  inches,  $9 \times 12$ , or "poster size."

Often people have a problem in mapping products into money. For simple choices, of course, such mappings are trivial. If a Snickers bar costs one dollar, you can easily figure out how much it costs to have a Snickers bar every day. But do you know how much it costs you to use your credit card? Among the fees you may be paying are: (a) an annual fee for the privilege of using the card (common for cards that provide benefits such as frequent flyer miles); (b) an interest rate for borrowing money (that depends on your deemed credit worthiness); (c) a fee for making a payment late (and you may end up making more late payments than you anticipate); (d) interest on purchases made during the month that is normally not charged if your balance is paid off but begins if you make your payment one day late; and (e) a charge for buying things in currencies other than dollars.

Credit cards are not alone in having complex pricing schemes that are neither transparent nor comprehensible to consumers. Think about mortgages, cell phone calling plans, and auto insurance policies, just to name a few. For these and related domains, we propose a very mild form of government regulation, a species of libertarian paternalism that we call RECAP: Record, Evaluate, and Compare Alternative Prices.

Here is how RECAP would work in the cell phone market. The government would not regulate how *much* issuers could charge for services, but it would regulate their disclosure practices. The central goal would be to inform customers of every kind of fee that currently exists. This would not be done by printing a long unintelligible document in fine print. Instead, issuers would be required to make public their fee schedule in a spreadsheetlike format that would include all relevant formulas. Suppose you are in Toronto and your cell phone rings. How much is it going to cost you to answer it? What if you download some email? All these prices would be embedded in the formulas. This is the price disclosure part of the regulation.

The usage disclosure requirement would be that once a year, issuers would have to send their customers a complete listing of all the ways they



had used the phone and all the fees that had been incurred. This report would be sent two ways, by mail and, more important, electronically. The electronic version would also be stored and downloadable on a secure Web site.

Producing the RECAP reports would cost cell phone carriers very little, but the reports would be extremely useful for customers who want to compare the pricing plans of cell phone providers, especially after they had received their first annual statement. Private Web sites similar to existing travel sites would emerge to allow an easy way to compare services. With just a few quick clicks, a shopper would easily be able to import her usage data from the past year and find out how much various carriers would have charged, given her usage patterns.\* Consumers who are new to the product (getting a cell phone for the first time, for example) would have to guess usage information for various categories, but the following year they could take full advantage of the system's capabilities. We will see that in many domains, from mortgages and credit cards to energy use to Medicare, a RECAP program could greatly improve people's ability to make good choices.

### Structure Complex Choices

People adopt different strategies for making choices depending on the size and complexity of the available options. When we face a small number of well-understood alternatives, we tend to examine all the attributes of all the alternatives and then make trade-offs when necessary. But when the choice set gets large, we must use alternative strategies, and these can get us into trouble.

Consider, for example, Jane, who has just been offered a job at a company located in large city far from where she is living now. Compare two choices she faces: which office to select and which apartment to rent. Suppose Jane is offered a choice of three available offices in her workplace. A

\*We are aware, of course, that behavior depends on prices. If my current cell phone provider charges me a lot to make calls in Canada and I react by not making such calls, I will not be able to judge the full value of an alternative plan with cheap calling in Canada. But where past usage is a good predictor of future usage, a RECAP plan would be very helpful.

reasonable strategy for her to follow would be to look at all three offices, note the ways they differ, and then make some decisions about the importance of such attributes as size, view, neighbors, and distance to the nearest rest room. This is described in the choice literature as a "compensatory" strategy, since a high value for one attribute (big office) can compensate for a low value for another (loud neighbor).

Obviously, the same strategy cannot be used to pick an apartment. In a large city like Los Angeles, thousands of apartments are available. If Jane ever wants to start working, she will not be able to visit each apartment and evaluate them all. Instead, she is likely to simplify the task in some way. One strategy to use is what Amos Tversky (1972) called "elimination by aspects." Someone using this strategy first decides what aspect is most important (say, commuting distance), establishes a cutoff level (say, no more than a thirty-minute commute), then eliminates all the alternatives that do not come up to this standard. The process is repeated, attribute by attribute (no more than \$1,500 per month; at least two bedrooms; dogs permitted), until either a choice is made or the set is narrowed down enough to switch over to a compensatory evaluation of the "finalists."

When people are using a simplifying strategy of this kind, alternatives that do not meet the minimum cutoff scores may be eliminated even if they are fabulous on all other dimensions. So, for example, an apartment that is a thirty-five-minute commute will not be considered even if it has a dynamite view and costs two hundred dollars a month less than any of the alternatives.

Social science research reveals that as the choices become more numerous and/or vary on more dimensions, people are more likely to adopt simplifying strategies. The implications for choice architecture are related. As alternatives become more numerous and more complex, choice architects have more to think about and more work to do, and are much more likely to influence choices (for better or for worse). For an ice cream shop with three flavors, any menu listing those flavors in any order will do just fine, and effects on choices (such as order effects) are likely to be minor because people know what they like. As choices become more numerous, though, good choice architecture will provide structure, and structure will affect outcomes.

Consider the example of a paint store. Even ignoring the possibility of

special orders, paint companies sell more than two thousand colors that you can apply to the walls in your home. It is possible to think of many ways of structuring how those paint colors are offered to the customer. Imagine, for example, that the paint colors were listed alphabetically. Artic White might be followed by Azure Blue, and so forth. While alphabetical order is a satisfactory way to organize a dictionary (at least if you have a guess as to how a word is spelled), it is a lousy way to organize a paint store.

Instead, paint stores have long used something like a paint wheel, with color samples ordered by similarity: all the blues are together, next to the greens, and the reds are located near the oranges, and so forth. The problem of selection is made considerably easier by the fact that people can see the actual colors, especially since the names of the paints are spectacularly uninformative. (On the Benjamin Moore Paints Web site, three similar shades of beige are called “Roasted Sesame Seed,” “Oklahoma Wheat,” and “Kansas Grain.”)

Thanks to modern computer technology and the World Wide Web, many problems of consumer choice have been made simpler. The Benjamin Moore Paints Web site not only allows the consumer to browse through dozens of shades of beige, but it also permits the consumer to see (within the limitations of the computer monitor) how a particular shade will work on the walls with the ceiling painted in a complementary color. And the variety of paint colors is small compared to the number of books sold by Amazon (millions) or Web pages covered by Google (billions). Many companies such as Netflix, the mail-order DVD rental company, succeed in part because of immensely helpful choice architecture. Customers looking for a movie to rent can easily search movies by actor, director, genre, and more, and if they rate the movies they have watched, they can also get recommendations based on the preferences of other movie lovers with similar tastes, a method called “collaborative filtering.” You use the judgments of other people who share your tastes to filter through the vast number of books or movies available in order to increase the likelihood of picking one you like. Collaborative filtering is an effort to solve a problem of choice architecture. If you know what people like you tend to like, you might well be comfortable in selecting products you don’t know, because people like you tend to like them. For many of us, collaborative filtering is making difficult choices easier.

A cautionary note: surprise and serendipity can be fun for people, and good for them too, and it may not be entirely wonderful if our primary source of information is about what people like us like. Sometimes it’s good to learn what people *unlike* us like—and to see whether we might even like that. If you like the mystery writer Robert B. Parker (and we agree that he’s great), collaborative filtering will probably direct you to other mystery writers (we suggest trying Lee Child, by the way), but why not try a little Joyce Carol Oates, or maybe even Henry James? If you’re a Democrat, and you like books that fit your predilections, you might want to see what Republicans think; no party can possibly have a monopoly on wisdom. Public-spirited choice architects—those who run the daily newspaper, for example—know that it’s good to nudge people in directions that they might not have specifically chosen in advance. Structuring choice sometimes means helping people to learn, so they can later make better choices on their own.<sup>5</sup>

### Incentives

Our last topic is the one with which most economists would have started: prices and incentives. Though we have been stressing factors that are often neglected by traditional economic theory, we do not intend to suggest that standard economic forces are unimportant. This is as good a point as any to state for the record that we believe in supply and demand. If the price of a product goes up, suppliers will usually produce more of it and consumers will usually want less of it. So choice architects must think about incentives when they design a system. Sensible architects will put the right incentives on the right people. One way to start to think about incentives is to ask four questions about a particular choice architecture:

Who uses?

Who chooses?

Who pays?

Who profits?

Free markets often solve all of the key problems by giving people an incentive to make good products and to sell them at the right price. If the market for sneakers is working well, there will be a lot of competition; bad

sneakers will be driven from the market and the good ones will be priced in accordance with people's tastes. Sneaker producers and sneaker purchasers have the right incentives. But sometimes incentive conflicts arise. Consider a simple case. When we go for our weekly lunch, each of us chooses his own meal and pays for what he eats. The restaurant serves us our food and keeps our money. No conflicts here. Now suppose we decide to take turns paying for lunch. Sunstein now has an incentive to order something more expensive on the weeks that Thaler is paying, and vice versa. (In this case, though, friendship introduces a complication; one of us may well order something cheaper if he knows that the other is paying. Sentimental but true.)

Many markets (and choice architecture systems) are replete with incentive conflicts. Perhaps the most notorious is the U.S. health care system. The patient receives the health care services that are chosen by his physician and paid for by the insurance company, with everyone from equipment manufacturers to drug companies to malpractice lawyers taking a piece of the action. Those with different pieces have different incentives, and the results may not be ideal for either patients or doctors. Of course, this point is obvious to anyone who thinks about these problems. But as usual, it is possible to elaborate and enrich the standard analysis by remembering that the agents in the economy are Humans. To be sure, even mindless Humans demand less when they notice that the price has gone up. But will they notice? Only if they are really paying attention.

The most important modification that must be made to a standard analysis of incentives is salience. Do the choosers actually notice the incentives they face? In free markets, the answer is usually yes, but in important cases the answer is no. Consider the example of members of an urban family deciding whether to buy a car. Suppose their choices are to take taxis and public transportation or to spend ten thousand dollars to buy a used car, which they can park on the street in front of their home. The only salient costs of owning this car will be the weekly stops at the gas station, occasional repair bills, and a yearly insurance bill. The opportunity cost of the ten thousand dollars is likely to be neglected. (In other words, once they purchase the car, they tend to forget about the ten thousand dollars and stop treating it as money that could have been spent on something else.) In contrast, every time the family uses a taxi the cost will be in their

face, with the meter clicking every few blocks. So a behavioral analysis of the incentives of car ownership will predict that people will underweight the opportunity costs of car ownership, and possibly other less salient aspects such as depreciation, and may overweight the very salient costs of using a taxi.\* An analysis of choice architecture systems must make similar adjustments.

Of course, salience can be manipulated, and good choice architects can take steps to direct people's attention to incentives. The telephones at the INSEAD School of Business in France are programmed to display the running costs of long-distance phone calls. If we want to protect the environment and to increase energy independence, similar strategies could be used to make costs more salient. Suppose the thermostat in your home was programmed to tell you the cost per hour of lowering the temperature a few degrees during the heat wave. This would probably have more effect on your behavior than quietly raising the price of electricity, a change that will be experienced only at the end of the month when the bill comes. Suppose in this light that government wants to increase energy conservation. Increases in the price of electricity will surely have an effect; making the increases salient will have a greater effect. Cost-disclosing thermostats might have a greater impact than (modest) price increases designed to decrease use of electricity.

In some domains, people may want the salience of gains and losses treated asymmetrically. For example, no one would want to go to a health club that charged its users on a "per step" basis on the Stairmaster. However, many Stairmaster users enjoy watching the "calories burned" meter while they work out (especially since those meters seem to give generous estimates of calories actually burned). Even better, for some, might be a pictorial display that indicated the calories one had burned in terms of food: after ten minutes one had earned only a bag of carrots but after forty minutes a large cookie.

We have sketched six principles of good choice architecture. As a concession to the bounded memory of our readers, we thought it might be

\*Companies such as Zipcar that specialize in short-term rentals could profitably benefit by helping people solve these mental accounting problems.

useful to offer a mnemonic device to help recall the six principles. By rearranging the order, and using one small fudge, the following emerges.

iNcentives

Understand mappings

Defaults

Give feedback

Expect error

Structure complex choices

Voilà: **NUDGES**

With an eye on these **NUDGES**, choice architects can improve the outcomes for their Human users.

## PART

# II

---

## MONEY

Not surprisingly, Humans differ dramatically from Econs in how they deal with money. Econs are sensible spenders and savers. They put money away for a rainy day, and for retirement, and they invest that money as if they had MBAs. When they borrow, Econs have no trouble choosing between fixed- and variable-rate mortgages, and they pay their credit card bills on time every month. If you are an Econ, you can skip this section of the book, unless you want to understand the behavior of your spouse, kids, and other Humans. A major goal of the next four chapters is to explore how people can do a better job at the difficult tasks of saving, investing, and borrowing. We also offer some suggestions about how private and public institutions might nudge people in directions that will make them a bit wealthier and more secure.