
DEFAULT

finding

How to choose a default

John Beshears, Richard T. Mason, & Shlomo Benartzi

abstract

We have developed a model for setting a default when a population is choosing among ordered choices—that is, ones listed in ascending or descending order. A company, for instance, might want to set a default contribution rate that will increase employees' average contributions to a retirement savings plan. A key input of the model is the distribution of *latent* options—the percentages of a population that select each available choice in the absence of a preset default. The model treats the default as an attraction point that causes some people to shift from their latent preference toward the default. It specifies the strength of each possible default's pull on each latent option and thereby points policymakers to the default most likely to achieve a desired aim. We tested our model using data from field experiments relating to retirement savings. In addition to presenting the results, which support the model's validity, we discuss how the model relates to prior empirical evidence on defaults.

Beshears, J., Mason, R. T., & Benartzi, S. (2022). How to choose a default. *Behavioral Science & Policy*, *8*(1), 1–15.

Core Findings

What is the issue?
Changing the default option—the option that is implemented on behalf of individuals who do not actively select an option for themselves—has been a major success for behavioral scientists seeking to influence policy-relevant outcomes. Choosing defaults, however, is difficult because the effects of defaults are nuanced. We provide guidance for choosing a default among ordered options when the goal is to increase the mean outcome in a population.

How can you act?
Selected recommendations include:
1) Examining experimental or nonexperimental data, as available, to determine the strength of defaults
2) When defaults are weak, placing a default just above a cluster of options that many individuals would select if forced to make a choice
3) When defaults are strong, placing a default well above such a cluster of options

Who should take the lead?
Policymakers and organizational leaders

When asked to identify the greatest contribution of behavioral economics to policy, Richard H. Thaler, winner of the Nobel Memorial Prize in Economic Sciences for his work on behavioral economics, pointed to improvements in retirement savings policies. One of the improvements involved changing the default from nonparticipation to participation in defined contribution plans, in which employees set aside a given amount for the plan from each paycheck.[1] Behavioral science research has documented that participation in such retirement plans increases dramatically when companies switch from an opt-in plan (in which no contribution is made unless an employee actively selects a contribution rate) to an opt-out plan (in which a preset, or *default*, contribution rate is used unless an employee actively changes the rate).[2–8] In a 2016 survey, 60% of 401(k) plans indicated that they use opt-out policies,[9] and such policies have been implemented at the national level in the United Kingdom, New Zealand, and Turkey.

The use of defaults is not limited to retirement plans. Changing the default changes the actions people take in domains such as organ donation,[10–12] insurance,[13] online marketing,[14] consumer product choice,[15] energy use,[16] tipping,[17] medication prescriptions,[18,19] and charitable donations.[20,21] A recent meta-analysis, combining data from multiple studies, documents the effects of defaults across settings.[22]

In this article, we offer guidance to policymakers who must choose a default from among many ordered options (see note A)—that is, ones listed in rank order, such as the percentage of salary contributed to a retirement plan. We present a model for determining whether and by how much any given default will cause individuals to deviate from their *latent* choice (the one that would be selected in the absence of a default) and for deciding which default to set for the greatest good of a population.

To aid readers in understanding the model, we speak mostly in terms of the concrete example of how it would be used for increasing contributions to defined contribution retirement plans. After describing the model, we report on field studies that illustrate how the model applies in a particular setting and provide empirical support for the model's predictions.

Model Basics

In applying the model to retirement plans, we take the perspective of a company policymaker who wishes to shift employee contribution rates upward (see note B). The policymaker may believe, for example, that a group of individuals is saving too little.[8] The insights that we develop in the context of retirement plans could help address challenges in other domains, such as encouraging a group of workers to lower their chosen thermostat settings in the winter[16] or inducing a group of physicians who are overprescribing brand-name medications to switch to prescribing generic equivalents.[18,19]

In a nutshell, we assume that a default is an attraction point. Any given default might cause some percentage of people to shift from their latent preference toward the default. When applying the model to a defined contribution retirement plan, we first collect data describing the distribution of latent contribution rates: For a subset of people in the population of interest, we observe the rate choices made in the absence of a default, and we identify the percentages who chose each latent value. Then we use the model to calculate the percentages of people who will shift from their latent value and how far they will shift in response to each of several defaults, with the resulting distributions of choices varying depending on different assumptions about the strength of the default's pull on nearby and distant latent values. The strength of the default is modeled by a combination of factors, or *parameters*, each of which can take any value in a range of values.

We then identify the combination of parameters that generates the model's most accurate predictions for how the distribution of contribution rates changes when a default is introduced. For a given combination of parameters, we compare the model-based predictions with actual choices made by a group of employees who demographically resemble the group that saw no default but were presented with a default. By finding the model-based predictions that most closely match the real-life choices that people make in response to several defaults, we identify the parameter values that best describe the population's responsiveness to defaults. Using those parameter values, we calculate the default that would be most effective at shifting the population's mean selection in a desired direction. (The model can also be useful when a policymaker's ability to collect data is limited. We address this situation in the article's last section.)

In principle, a policymaker seeking to increase the mean contribution rate of a group of employees should select a default that pulls many contribution rates up while pulling few down. Our modeling indicates that to achieve this aim, a policymaker should choose a default that is above a cluster of popular latent contribution rates. When the default has a weak influence, the default that maximizes the mean contribution rate is likely to be one that is only slightly above that cluster. When the default has a strong influence, policymakers will maximize the mean contribution rate by setting a default that is significantly above the cluster. To more fully explain how the model guides policymakers to a desirable default, we provide concrete details in the next section and in the Supplemental Material. We should note that we consider our model to be empirical rather than explanatory, because it does not specify the economic and psychological mechanisms driving people's responses to defaults; it merely describes how different defaults shape the distribution of the selections made by a population. As we explain later, however, its assumptions are consistent with the findings of past behavioral science research into defaults.

The Model in Detail

Our model has three ingredients: (a) the distribution of latent contribution rates, (b) the value of a default, and (c) a set of formulas that uses three parameters to predict whether and to what extent the default will affect an individual who has a particular latent contribution rate.

In applying the model to retirement savings plans, we obtained the distribution of latent

"In a nutshell, we assume that a default is an attraction point"

choices from data we collected in one of three empirical studies we conducted (these are described more fully later in the article). In this study—which we call Study 1 despite its not having been done first—one group of employees visited a website to select their contribution rate to a company's retirement program and saw no default; they used a keyboard to enter a number into a blank space. After entering a number, they could either retain that initial number or select a different number, whether by interacting with the website during the same session or by returning to the website at a later time. We defined an individual's latent contribution rate (L) as the individual's selected rate at the end of this process. Other groups of employees in Study 1, as well as employees in our two other empirical studies (Studies 2 and 3), saw a prepopulated number where employees in the first group saw a blank space. Otherwise, their experience was identical to that of employees in the first group. The prepopulated number is the default value (D), and in our studies, it took values in the range of 6% to 11%.

The third ingredient—the heart of the model—uses two formulas to govern individuals' responses to defaults, with the three parameters mentioned above determining how strongly defaults influence choices.

The first formula predicts the likelihood that an individual with a given latent choice will end up changing their choice in response to a given default. It incorporates a parameter termed the *radius* (R) and an *adjustment factor* (F).

The R value indicates how far a default's influence extends. If the default were a point on a horizontal line indicating the potential contribution rates in ascending order, the radius would describe the distance to the right and to the left within which the default has an effect on latent rates. For instance, an R of 4 around a default of 7% indicates that the default has an effect on

60%

401(k) plans that indicated use of opt-out policies in 2016

Heart of the model

2 formulas that represent individuals' responses to defaults

psychological anchoring

Tendency for a person asked to choose a numerical value to start with some reference point and then only slightly adjust away from it

individuals whose latent values are between 3% and 11%.

The formula reflects the presumption, based on past research into defaults, that a default's effect declines as the latent rate goes in either direction from the default. When a person's latent contribution rate is close to the default, the default is likely to cause the person to shift from the latent rate in the direction of the default. When a person's latent contribution rate is far from the default, the default is less likely to influence the outcome. In general, when the default is above an individual's latent contribution rate, the default is likely to pull the individual's contribution rate higher. Conversely, when the default is below the latent contribution rate, the default may pull down the final contribution rate selection.

In the model, an individual whose latent contribution rate is equal to the default is deemed to be 100% likely to be influenced by the default. An individual whose latent contribution rate is a distance of exactly R away from the default is deemed to have zero probability of being influenced by the default. The probability that the default has an effect declines linearly between those two points, except when the latent contribution rate is especially attractive on its own, an issue that we discuss next. (See note C for a caveat.)

The parameter F is an adjustment factor for the counterpull exerted by latent values that are especially attractive on their own. In retirement savings plans, our own research and other data sets have shown that people like to choose contribution rates that are multiples of five, and they tend to resist shifting to other values; that is, in common with defaults, multiples of five are attraction points. We adjust for this counterpull by decreasing the calculated pull of the default by the amount of resistance generated by latent multiples of five. In the model, we address the counterpull with a formula, or *function*, called $A(L)$. $A(L)$ is set to zero if the latent contribution rate L is not a multiple of five, because these latent rates are assumed to exert no resistance. For multiples of five, $A(L)$ is set to F, with F taking a value in the range from 0 to 1 ($0 \leq F \leq 1$).

In mathematical notation, we say that when $|D - L| \leq R$ (that is, when the difference between the default value D and the latent contribution rate L is less than or equal to the parameter R), the probability that the individual is influenced by the default is $[1 - A(L)] \times (1 - |D - L|/R)$. When $|D - L| > R$, the probability that the individual is influenced by the default is zero.

If an individual is influenced by the default, the individual ends up with a contribution rate calculated by the second formula at the heart of our model: $[(1 - W)L + WD]$. This formula yields a weighted average of the latent contribution rate and the default. The weighting factor (W) takes a value between 0 and 1 ($0 \leq W \leq 1$). For example, a W of 0.7 (giving fairly strong weight to the default) paired with a D of 10 and an L of 0 would result in a final contribution rate of 7, with the individual moving seven tenths of the way from their latent value toward the default. The online interface in our three studies encouraged individuals to choose contribution rates as whole number percentages, so we round the contribution rate given by the formula to the nearest whole number.

To illustrate how the model works, we have constructed two examples. In the first, depicted in Figure 1, the model uses the parameters R = 12, F = 0.3, and W = 0.9, a combination indicative of the default having a strong, far-reaching influence and multiples of five exerting a modest resistance. The white bars reflect the distribution of actual latent contribution rates revealed by participants in Study 1. The gray and black bars show the predicted distribution of contribution rates when the default was modeled at 7% or 10%, respectively. The predictions indicate that when a default exerts a strong effect over a wide range of latent values, a wide swath of the distribution will be drawn toward the default. Notably, regardless of people's latent rates, a high percentage of those who were modeled to have been presented with a default of 7% switched their choices, and the percentage who chose 7% rose from less than 5% to more than 50%.

The second example uses the same defaults and distribution of latent contribution rates as in the first example (see Figure 2). The defaults'

Figure 1. Illustrative model prediction of contribution rates when a default exerts a strong effect on rates (parameter values: R = 12, F = 0.3, W = 0.9)

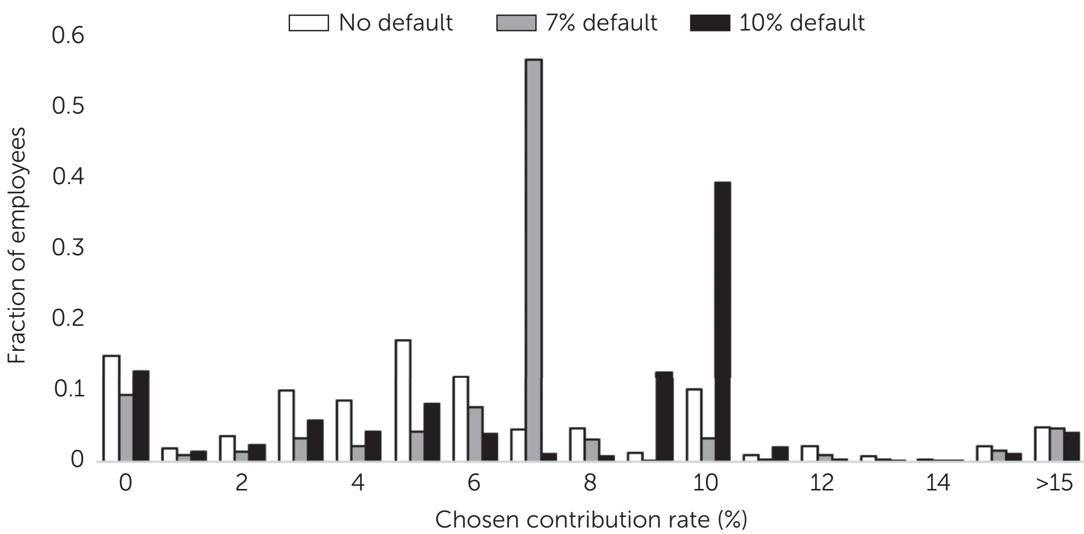

Note. The white bars show the distribution of contribution rates in the no default condition (that is, the latent rates); those rates are drawn from the empirical data in Study 1. The gray and black bars show the distribution of contribution rates predicted by the model for a 7% default and for a 10% default, respectively. The results indicate that when a default exerts a strong effect, people having latent preferences both near to and far from the default will be drawn toward the default—as is evidenced by the declines in the fraction of employees choosing many of the rates. See the main text for definitions of *R*, *F*, and *W*.

Figure 2. Illustrative model prediction of contribution rates when a default exerts a weak effect on rates (parameter values: R = 1.5, F = 0.3, W = 0.7)

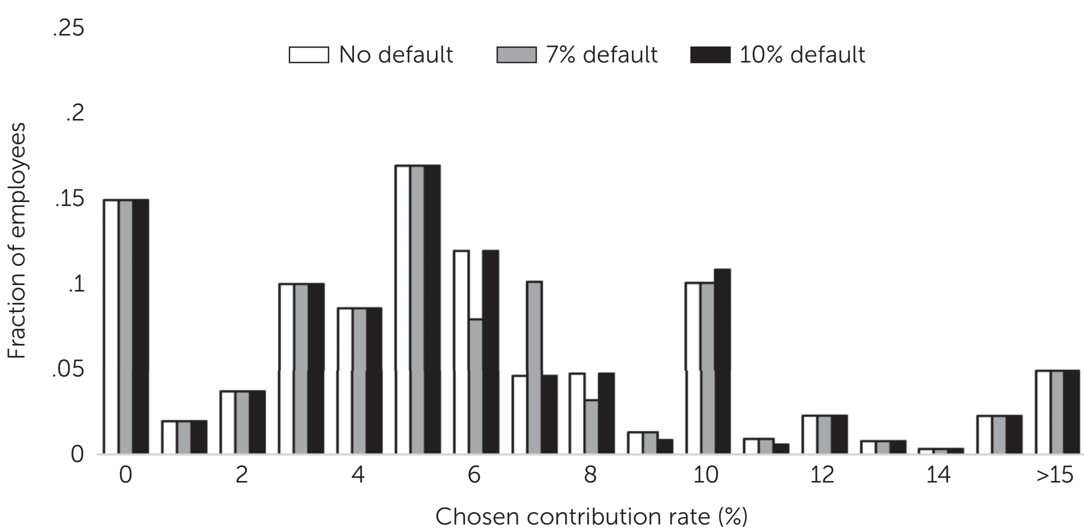

Note. As in Figure 1, the white bars show the distribution of contribution rates in the no default condition (that is, the latent rates), and the gray and black bars show the distribution of contribution rates predicted by the model for a 7% default and for a 10% default, respectively. The results indicate that even when the 7% default has a weak effect, it nonetheless exerts a draw on individuals whose latent contribution rates are 6% or 8% (as indicated by the decline in the fraction of employees predicted to choose those rates). The 10% default has less of an effect on the distribution. The parameter values used here generate predictions that most closely match (that is, are the best fit for) the empirical findings from Studies 1, 2, and 3—a correspondence implying that the parameter combination is the best for predicting the responses to defaults in a real-life population resembling that in our studies.

effects are assumed to be weaker, however, as is reflected in the parameters $R = 1.5$, $F = 0.3$, and $W = 0.7$. The output implies that despite the weak effect of the 7% default (gray bars), this default still exerts a draw on individuals whose latent contribution rates are 6% or 8%. The 10% default (black bars) has less of an impact on the distribution. It exerts a pull on individuals whose latent contribution rates are 9% or 11%, but such an influence is less meaningful because few individuals have those latent contribution rates.

This last combination of parameters yielded the best fit with our experimental data; that is, it most closely replicated the outcomes we found when the real-life employees we studied were presented with a default of 6%, 7%, 8%, 9%, 10%, or 11%. Later in the article, we address the implications of this finding for setting defaults, but first we describe the empirical studies we conducted and the ways in which they support the validity of our model and confirm past research on defaults.

The Experiments

Experimental Design

We conducted three experiments, all of which were completed before we conducted our modeling. Even though the experiment we call Study 1 was not run first, we treat it as our primary study because it was the only one that enabled us to observe the distribution of latent contribution rates and thus to study the effects of various defaults on that distribution.[20,21,23–25] We describe Study 1 in this subsection and address Studies 2 and 3, which were similar, in a later subsection. For details, see the Supplemental Material.

We worked with the segment of the company Voya Financial that helps employers manage retirement savings plans. For a subset of employers, employees who became eligible for the retirement plan were invited to visit a Voya-administered website, Voya Enroll, to begin contributing. Figures S1–S8 in the Supplemental Material show screenshots of what employees saw during the registration process.

In Study 1, we assigned employees randomly to one of three groups when they reached the webpage at which they selected their contribution rates. The study had three conditions: 7% default, 10% default, and no default.

As shown in Table 1, approximately half of the individuals in Study 1 were men. The mean age of participants was 38 years, and their mean annual salary was approximately $70,000. These characteristics did not show statistically significant differences across the three conditions.

In the 7% default and 10% default conditions of Study 1, the space for indicating the desired contribution rate was prepopulated with the

Table 1. Participant gender, age, & salary in Study 1, by randomly assigned condition

	Experimental condition			
Characteristic	No default	7% default	10% default	*p* value from χ^2 or *F* test for null hypothesis that conditions are equal
% men	53	52	52	.66
Age in years				.69
Mean	38	38	38	
Standard deviation	12	12	12	
Salary				.16
Mean	$69,000	$71,000	$71,000	
Standard deviation	$51,000	$52,000	$54,000	
Number of participants	3,991	4,024	4,048	

Note. The *p* values indicate that the three participant groups do not differ significantly. *Standard deviation* is a measure of the amount of variation in a set of values; approximately two thirds of the observations fall between one standard deviation below the mean and one standard deviation above the mean.

default of interest (see Figure S4 in the Supplemental Material). In the no default condition, the space for indicating the desired contribution rate was empty when the webpage loaded, and a blinking cursor prompted the employee to enter a number. (See Figure S5 in the Supplemental Material.) As soon as a number was entered, the webpage transformed to appear as if the entered number had been the prepopulated contribution rate (as in Figure S4 in the Supplemental Material). In all three conditions in Study 1, employees could increase or decrease their chosen contribution rate away from the initial rate by clicking on the + or – keys.

As specified when we preregistered Study 1 (see Figure S9 in the Supplemental Material), our primary outcome variable is the contribution rate in effect 60 days after the initial Voya Enroll visit, adjusted to reduce the potentially misleading influence of outliers by setting values below the 1st percentile equal to the 1st percentile and values above the 99th percentile equal to the 99th percentile. Preregistration is done for transparency, that is, to minimize the likelihood that researchers will cherry-pick data and thus publish misleading results. The choice of a 60-day window balances two factors. On the one hand, a longer time window would increase the likelihood that factors unrelated to the default, such as salary increases or financial emergencies, could influence the final contribution rates. On the other hand, a shorter time window might miss changes that employees make after having some time to ponder their choice more fully. Some employees choose not to enroll in the plan when they first visit Voya Enroll but return within a few weeks and select a positive contribution rate.

Results

As we pledged in our preregistered analysis plan for Study 1, we calculated, using the analytic method known as *ordinary least squares regression*, the effect of the 7% or 10% default on the mean contribution rate. Relative to having no default, the 7% default decreased the mean contribution rate by 0.02 percentage points when we did not control for gender, age, and salary and by 0.04 percentage points when we accounted for those factors; the 10% default

“defaults can trigger shifts from latent values”

increased the mean contribution rate by 0.08 percentage points when we omitted controls and by 0.06 percentage points when we included controls. None of these estimates were statistically significant, and all of them were small in magnitude. When we used the same analytic approach to investigate whether the 7% default and 10% default increased the likelihood that an individual would choose a contribution rate greater than zero, we similarly found that the effects were not statistically significant and were small in magnitude. We had hypothesized that the 7% default and 10% default would increase a population’s mean contribution rate relative to having no default, so we were surprised by these results.

Such findings could have implied that setting defaults did not influence contribution decisions, but further analyses, which were not preregistered, indicated that defaults did, indeed, affect contribution decisions, even though they did not affect mean contribution rates. The results support the idea that defaults can trigger shifts from latent values among people who are signing up for retirement plans even when the average rate for the population does not change in a desired direction.

The data from Study 1 also showed that study participants were attracted to contribution rates that were multiples of five, as previous work has found.[4] This attraction is evident in Figure 3, which shows the distribution of the final contribution rates in the three experimental conditions. It is because of this finding that our model assumes that individuals whose latent contribution rates are multiples of five are less likely than others to be influenced by defaults. (See note D.)

Additional analyses revealed specific influences of defaults, including that the defaults in the study increased the fraction of individuals who ended up with contribution rates equal to the default. To identify this pattern, we compared

Figure 3. Distributions of contribution rates by randomly assigned condition in Study 1

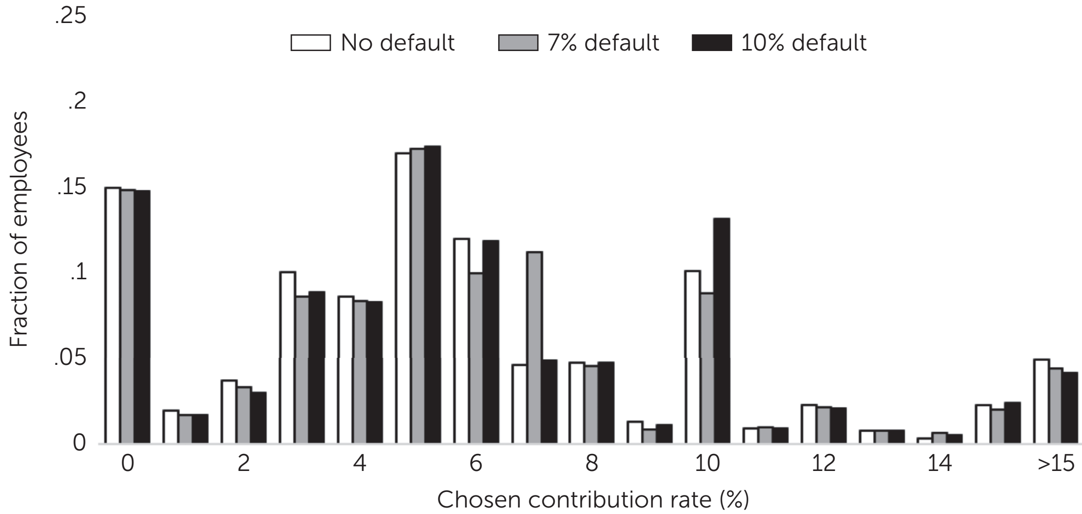

Note. The white, gray, and black bars show the distribution of contribution rates in Study 1 in the no default condition, the 7% default condition, and the 10% default condition, respectively. The data reflect a relatively weak influence of the defaults but show some shifting of latent values toward the defaults—as is indicated by increases in the fraction of employees selecting the defaults and declines in the fractions choosing several other rates.

the fraction of employees who chose a given contribution rate (termed *C*%, with *C* being an integer) in the 7% default and 10% default conditions with the fraction of employees who chose that rate in the no default condition. Using ordinary least squares regression again, we calculated the differences separately for each integer contribution rate from 0% to 15%. Because very few people select higher contribution rates, we treated all contribution rates equal to 16% or higher as belonging to a single category. In Figure 4, the contribution rate varies along the horizontal axis. The light and dark vertical bars indicate, respectively, the effect of the 7% default or the 10% default on the likelihood of a given contribution rate being chosen, relative to the likelihood when no default was presented. This effect is measured in terms of the size of the difference in the percentage of employees who chose the given contribution rate. The I-shaped lines, commonly known as *whiskers*, give 95% confidence intervals; the findings are statistically significant when the whiskers do not pass through the horizontal zero line. (See note E for a definition of 95% confidence intervals.)

The data indicate that relative to having no default, the 7% default caused a statistically significant increase in the fraction of individuals with a 7% contribution rate, and the 10% default caused a statistically significant increase in the fraction of individuals with a 10% contribution rate. This finding, too, supports our model, which predicts that some individuals with latent contribution rates close to a default will end up choosing the default. This finding is also consistent with prior literature documenting that defaults are chosen frequently.[2–6,17,20,21]

The analyses also revealed—again consistent with our modeling and with previous findings[5,6,20,21]—that individuals sometimes ended up choosing the default either when their latent contribution rate was below the default or when their latent contribution rate was above the default. The 7% default decreased the fraction of employees with a contribution rate less than or equal to 6% and decreased the fraction of employees with a contribution rate greater than or equal to 8%. The 10% default decreased the fraction of employees with a contribution rate less than or equal to 9% and decreased the fraction of employees with a contribution rate greater than or equal to 11%, although the last

Figure 4. The effect of the 7% default & the 10% default on the likelihood of the contribution rate being equal to a given value in Study 1

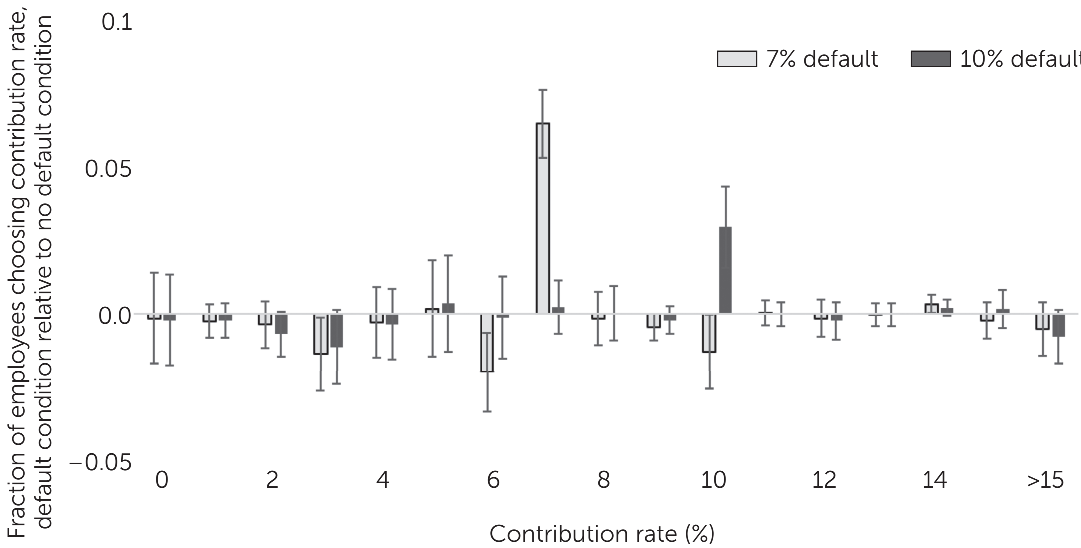

Note. The plot here compares the fraction of employees who chose a given contribution rate in the 7% default or 10% default condition with the fraction of employees who chose that rate in the no default condition. The results indicate that the defaults increased the fraction of individuals who ended up with contribution rates equal to the default. The data support our model, which predicts that some individuals with latent contribution rates close to a default will end up choosing the default. The whiskers show 95% confidence intervals; findings are statistically significant when the whiskers do not pass through the horizontal zero line.

finding was not statistically significant, perhaps because the fraction of employees with a latent contribution rate greater than or equal to 11% is so low that there is little room to decrease it further.

Finally, in line with our modeling, we found some evidence consistent with past research that indicated a default is more likely to influence an individual whose latent contribution rate is close to the default than an individual whose latent contribution rate is far from the default[5,6,17,20,21]—although the statistical power of the tests we did is low. (See note F. Also see the Additional Statistical Tests section of the Supplemental Material for more detail about the analyses relating to this finding.)

Calibrating the Model

To see if the model could be useful for indicating which default would be best for raising the mean contribution rate of a given population, we used the following process. To *calibrate the model*, we examined the effects of all possible combinations of *R*, *F*, and *W* on the distributions of contribution rates in the presence of default rates and looked for the distributions that best matched the real-world distributions found in our empirical studies. The combination of parameter values that led to the best-fit distributions could be presumed to predict the behavior of other populations whose demographic characteristics were similar to those of the study participants when those populations encountered retirement sign-up programs similar to those our participants encountered.

To find the best fit, we combined data from Study 1 with data from Studies 2 and 3. Studies 2 and 3 were conducted prior to Study 1 and were not preregistered. They had the same design as Study 1 except that they lacked a no default condition and had conditions with integer default contribution rates of 6% through 11%, rather than solely 7% and 10%. See the Supplemental Material for more details.

Given the distribution of latent contribution rates from the no default condition of Study 1, we input all possible combinations of the parameters into the model. *R* ranged from 0.5 to 15.0 in increments of 0.5; *F* ranged from 0 through 1.0

in increments of 0.1; and *W* ranged from 0 to 1.0 in increments of 0.1. Using each combination, the model predicted the distribution of contribution rates when employees were presented with a 6%, 7%, 8%, 9%, 10%, or 11% default. We then compared the model's predictions with the observed distributions of contribution rates of the participants, and we calculated how closely the modeled distributions matched those found in the empirical studies. See the Supplemental Material for more details.

As we mentioned earlier, for the default values examined in our studies (6%, 7%, 8%, 9%, 10%, and 11%), the model best fits the data when *R* takes a value of 1.5, *F* takes a value of 0.3, and *W* takes any value in the interval $0.5 < W \leq 1.0$ (with $R = 1.5$, the model makes the same predictions for all of these values of *W*).

Of course, the model's predictions using the best-fitting parameter values do not capture every feature of the real-world data. For example, the model with these parameter values predicts that the default does not affect individuals whose latent contributions rates are two percentage points or more away from the default. However, in the data from Study 1, the 7% default condition leads to a statistically significant 2.2 percentage point decrease in the fraction of individuals who choose contribution rates of 5% or less and a statistically significant 2.4 percentage point decrease in the fraction of individuals who chose contribution rates of 9% or more, relative to the no default condition (because those individuals moved toward the default). Similarly, the 10% default condition leads to a statistically significant 2.2 percentage point decrease in the fraction of individuals who chose contribution rates of 8% or less, relative to the no default condition. (The 10% default condition did not have a statistically significant effect on the fraction of individuals who chose contribution rates of 12% or more, relative to the no default condition.) Overall, though, the best-fitting parameter values for the model include a low value of *R*, indicating that whatever the default rate is, it tends to attract individuals whose latent contribution rates are close to the default.

Figure 5. Mean contribution rate predicted by the model with best-fit parameter values ($R = 1.5$, $F = 0.3$, $W = 0.7$) as the default varies

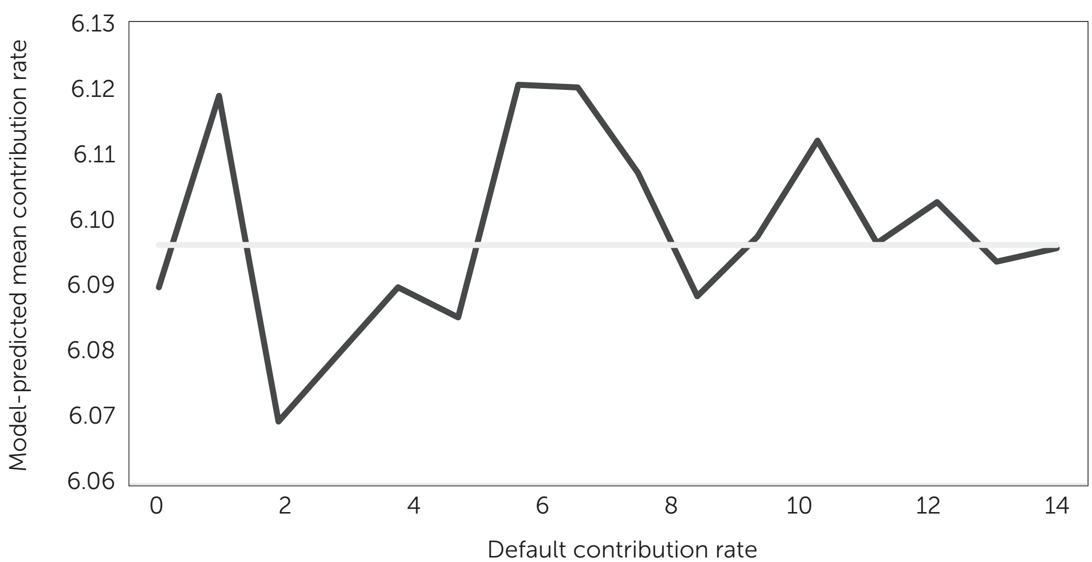

Note. Using the model parameters that produced contribution-rate distributions most like those in the empirical studies ($R = 1.5$, $F = 0.3$, $W = 0.7$), we determined that setting a default of 6% or 7% would result in the highest mean contribution rate in a population that resembled the one in our empirical studies. The horizontal line in the middle of the figure shows the mean contribution rate in the no default (that is, latent) condition. To arrive at the means shown, we calculated the model's predictions for the distribution of contribution rates in response to each possible integer default and then computed the mean of that distribution. One benefit of the model is that it makes predictions about contribution-rate distributions for defaults that we did not test in our experiments (defaults less than 6% or greater than 11%).

Figure 5 shows the model's predictions, given the best-fit parameter values, for the mean contribution rate as the default varies. The model-predicted mean reaches a peak at a default of 6%, and the mean for a default of 7% is nearly identical. In other words, to maximize the mean contribution rate of a population that resembles the one in the studies, the model indicates that policymakers would set a 6% or 7% default. (See note G.) The mean contribution rate for a population is calculated as follows: For each possible contribution rate, we multiply the contribution rate by the fraction of the population predicted by the model to choose that contribution rate, then we calculate the sum across contribution rates.

Comparisons With Other Models

As we mentioned at the start of this article, we constructed our model without specifying the mechanisms driving individuals' responses to defaults. Nonetheless, the model can be compared with ones that articulate mechanisms for the effects of defaults.

Our model implies that individuals whose latent contribution rates are closer to the default are more likely to be affected by the default, a feature consistent with models assuming that people incur a cost—in the form of inconvenience—if they opt out of a default.[21,23,26] In these models, individuals' latent contribution rates are assumed to be their most preferred contribution rates; therefore, individuals whose latent contribution rates are farthest from the default have the strongest incentive to bear the inconvenience of opting out and switching from the default to their most preferred contribution rates. Individuals whose latent contribution rates are close to the default have a weaker incentive to go to the trouble of opting out and are more likely to remain at the default.

Yet our model differs from those that focus on the costs of opting out in that, like models based on the phenomenon known as *psychological anchoring*,[26] it allows for the possibility that people who go to the trouble of rejecting the default will choose a contribution rate close to the default instead of choosing their latent

"individuals whose latent contribution rates are closer to the default are more likely to be affected by the default"

rate. *Anchoring* refers to the tendency for a person asked to choose a numerical value to start with some reference point and then only slightly adjust away from it. In models that focus on opt-out costs, the default does not attract individuals to contribution rates close to but not equal to the default, whereas models that focus on psychological anchoring allow that kind of attraction.

It would be desirable to determine whether the effects of defaults in our empirical setting are driven by opt-out costs or psychological anchoring, but the data do not allow us to distinguish between these two mechanisms. Models that view the default as an anchor allow the default to cause an increase in the fraction of individuals choosing a contribution rate near to but not equal to the default.[26] However, these models also feature a countervailing force: Individuals who have that nearby contribution rate as their latent contribution rate are likely to move from their latent contribution rate to the default. On net, the default can lead to either an increase or a decrease in the fraction of individuals choosing that nearby contribution rate. In our real-world data, the default decreases the fraction of individuals choosing nearby contribution rates, but this evidence cannot distinguish between a model of anchoring and a model of opt-out costs because both types of models can predict this empirical pattern.

Policy Implications

Our model applies to many contexts beyond retirement savings. The designer of a smart thermostat can set the default temperature that a home's heating and cooling system targets. The designer of an electronic health record system can set the default number of pills prescribed by a physician for a given patient profile and medication. The designer of a webpage for charitable contributions can set the default donation amount. The model parameter values that best fit our experimental data are unlikely to be the parameter values that are appropriate when applying the model in other domains. Nonetheless, as we discussed earlier in the article, evidence from a variety of contexts supports the assumptions of the model, suggesting that the structure of the model is indeed applicable in a range of settings.

If a policymaker is trying to increase the mean outcome for a population on some measure, the model provides guidance for selecting a default among ordered options. In general terms, the policymaker should first identify the distribution of latent outcomes. Next, the policymaker should gauge how influential the default is. This information, in turn, should be used to set a default that will pull up the outcomes of many individuals while pulling down the outcomes of few individuals. If the default is weak (that is, if the radius, *R*, within which the default has an effect, is small), the default that maximizes the mean outcome is likely just above a cluster of popular latent outcomes. If the default is strong (that is, *R* is large), the default will likely be higher. When *F* (the value of the adjustment factor for focal, or sticky, options, such as multiples of five) is high, the policymaker should generally avoid placing the default just above latent outcomes that individuals are reluctant to leave (because the default would then pull few individuals up) and should try to place the default just below such latent outcomes (because the default would then pull few individuals down).

For a policymaker to implement this guidance, the ideal approach would be to run an experiment similar to Study 1, featuring a condition with no default (to observe the distribution of latent outcomes) and conditions with defaults (to estimate the strength of the default). If this approach is not feasible, nonexperimental data can be informative. For example, if a company is using a given default at sign-up and finds that few individuals end up with the default option, program managers can infer that the influence of the default is weak and that the distribution of observed outcomes approximates the distribution of latent outcomes. If managers find that many individuals choose the default option, they can infer that the default's influence is strong. In this latter case, the policymaker would want to push the default to be more extreme so as to shift outcomes in the desired direction.

For additional insight into the likely strength of the default, a policymaker who cannot conduct a study can rely on past research. According to prior work,[22] defaults are more effective in domains where individuals are asked to make consumer purchase decisions and less effective in domains where individuals are asked to make pro-environmental decisions, and they are more influential when they communicate the policymaker's recommendation[27] or serve as a reference point against which other options are judged[28] than when they merely make the default option easy to implement.[23]

Our analysis has limitations. The model applies to many settings but not all. For example, in situations where the default influences outcomes primarily because many people are inattentive—that is, they do not notice that a default is being implemented—the model's assumptions regarding the way in which defaults influence outcomes may not be satisfied.[29] In these situations, it is less likely to be true that the influence of the default gets weaker as the difference between the default and an individual's latent outcome increases.

This observation highlights a key feature of our experimental setting. The participants made a choice to visit a website for enrolling in a retirement savings plan, so they were paying attention to the decision at hand.[21] This fact may explain why the default effects we observed are weaker than some other default effects that have been documented previously in studies of retirement savings plans.[2–6] Perhaps the individuals in our experiment arrived at the website having already thought through the contribution rate they would like to choose and were therefore less susceptible to the default's influence.

In this article, we have not addressed the moral considerations that a policymaker should have in mind when choosing a default. We have adopted the perspective of a policymaker who is trying to shift outcomes in a particular direction for ethically sound reasons. For example, the policymaker may have strong reasons to believe that psychological biases are causing individuals' choices to deviate systematically from the choices that would maximize their welfare. For another example, the policymaker may wish to shift outcomes because people are making decisions in ways that do not account for the consequences of their choices on others, such as when people consume energy excessively without regard for their contribution to global climate change. Policymakers who are unsure of which outcomes are appropriate should use a different framework for contemplating default selection.[26] They should also be careful to avoid subjecting individuals to the risk of significantly negative outcomes.

Our analysis points to some interesting extensions. We considered the choice of a single default for a population of individuals. If those individuals can be divided into easily identifiable subpopulations who have different latent distributions than the full population does and who respond differently to particular defaults (in other words, whose choices are described by different model parameters), it would be possible to tailor a different default for each subpopulation. This line of reasoning can be applied to situations in which the policymaker has a more complex objective than simply shifting mean outcomes of a large population upward or downward. For example, if a policymaker believes that individuals with low incomes have a greater or lesser need for higher retirement plan contribution rates than individuals with high incomes, default policies could be adjusted on the basis of income, with one group's default chosen to increase contribution rates and the other group's default chosen to promote more moderate contribution rates.

As another extension, it would be valuable to consider how a default might change over time. Consider the case of a smart thermostat. To reduce energy consumption at a company, the building managers might initially begin with a default temperature that is only slightly below the temperature that employees would choose for themselves during winter. After a set time, as the workers habituated to colder temperatures, the managers might lower the default.

Defaults affect the distribution of outcomes in subtle ways. By using our model, policymakers can select defaults for maximal impact.

endnotes

A. We do not address situations featuring a small number of options (say, five or fewer) in a choice menu. Our model could accommodate such situations, but the structure imposed by our model would be unnecessary. We also do not address situations featuring many unordered options, because our model does not speak to those situations.

B. The model could also be adjusted for use by policymakers who, for whatever reason, wanted to shift contribution rates downward.

C. Although the model puts the probability of a default's influence at zero if the latent value is beyond the default's radius of effect (that is, when $|D - L| > R$), in reality, even individuals whose latent contribution rates are very far from the default have some chance of being influenced by the default. We treat the probability as zero for simplicity, on the grounds that it is likely to be much closer to zero than is the case when the latent value is close to the default.

D. One could make the argument that people whose latent contribution rates are multiples of five might be more likely to be influenced by defaults because they have thought less deeply about their contribution rate choices. However, as we show in the Calibrating the Model section, our calibration exercise indicates that giving F a strictly positive value—0.3—gives the best fit for the data, suggesting that the assumption embedded in our model is the correct one. For additional evidence on the attractiveness of round numbers, see reference 30.

E. Editors' note to nonscientists: A 95% confidence interval for a given metric indicates that in 95% of random samples from a given population, the measured value will fall within the stated interval.

F. The model assumes that the effect of the default will be the same for latent values that are an equal distance below or over the default. Additional analyses that test this assumption of a symmetric effect around the default are described in the Additional Statistical Tests section of the Supplemental Material. The evidence does not contradict the assumption, but the statistical power of the test is low. We view the issue as an interesting one for future research to address.

G. We do not view these predictions as contradictory to the empirical data because the predicted values fall within the 95% confidence intervals of the corresponding empirical estimates.

author affiliation

Beshears: Harvard Business School and National Bureau of Economic Research. Mason: Voya Financial and Carnegie Mellon University. Benartzi: University of California, Los Angeles. Corresponding author's e-mail: jbeshears@hbs.edu.

author note

We thank Thomas Armstrong and Daniella Listro of Voya Financial for implementing the experiments and preparing the data, and we thank Katherine Milkman for her involvement in the early stages of this research. Andrew Joung, Seung Hyeong Lee, Predrag Pandiloski, Byron Perpetua, and Kartikeya Vira provided outstanding research assistance. We have benefited from the comments of Thomas Armstrong, Saurabh Bhargava, Lombard Gasbarro, Daniella Listro, Marilyn Morgan, Charles Nelson, Mark Patterson, John Payne, Stephen Shu, Stephen Thomas, and Giovanni Urga. Benartzi is a paid consultant to Voya Financial. Beshears and Benartzi have, at various times, received compensation from and sat on the advisory boards of financial firms. See their websites for complete lists of outside activities. Voya Financial had the opportunity to review the article before public release for the sake of identifying factual inaccuracies, but the authors retained full editorial control of the article contents. The findings and conclusions expressed are solely those of the authors and do not represent the views of Voya Financial, Harvard University, the National Bureau of Economic Research, Carnegie Mellon University, or University of California, Los Angeles.

supplemental material

- https://behavioralpolicy.org/publications/
- Method & Analysis

references

1. Dubner, S. J. (Host). (2018, December 19). People aren't dumb. The world is hard (No. 340) [Audio podcast episode]. In *Freakonomics radio.* Freakonomics Radio Network. https://freakonomics.com/podcast/people-arent-dumb-the-world-is-hard-rebroadcast
2. Madrian, B. C., & Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *The Quarterly Journal of Economics*, *116*(4), 1149–1187. https://doi.org/10.1162/003355301753265543
3. Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2002). Defined contribution pensions: Plan rules, participant decisions, and the path of least resistance. In J. M. Poterba (Ed.), *Tax policy and the economy* (Vol. 16, pp. 67–113). MIT Press.
4. Choi, J. J., Laibson, D., & Madrian, B. C. (2004). Plan design and 401(k) savings outcomes. *National Tax Journal*, *57*(2, Pt. 1), 275–298. https://doi.org/10.17310/ntj.2004.2.07
5. Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2004). For better or for worse: Default effects and 401(k) savings behavior. In D. A. Wise (Ed.), *Perspectives on the economics of aging* (pp. 81–121). University of Chicago Press. https://doi.org/10.7208/chicago/9780226903286.003.0003
6. Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (2008). The importance of default options for retirement saving outcomes: Evidence from the USA. In S. J. Kay & T. Sinha (Eds.), *Lessons from pension reform in the Americas* (pp. 59–87). Oxford University Press.
7. Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness.* Yale University Press.
8. Benartzi, S., & Thaler, R. H. (2013, March 8). Behavioral economics and the retirement savings crisis. *Science*, *339*(6124), 1152–1153. https://doi.org/10.1126/science.1231320
9. Plan Sponsor Council of America. (2018). *60th annual survey of profit sharing and 401(k) plans.*
10. Johnson, E. J., & Goldstein, D. (2003, November 21). Do defaults save lives? *Science*, *302*(5649), 1338–1339. https://doi.org/10.1126/science.1091721
11. Abadie, A., & Gay, S. (2006). The impact of presumed consent legislation on cadaveric organ donation: A cross-country study. *Journal of Health Economics*, *25*(4), 599–620. https://doi.org/10.1016/j.jhealeco.2006.01.003
12. Steffel, M., Williams, E. F., & Tannenbaum, D. (2019). Does changing defaults save lives? Effects of presumed consent organ donation policies. *Behavioral Science & Policy*, *5*(1), 69–88. https://doi.org/10.1353/bsp.2019.0005
13. Johnson, E. J., Hershey, J., Meszaros, J., & Kunreuther, H. (1993). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty*, *7*(1), 35–51. https://doi.org/10.1007/BF01065313
14. Johnson, E. J., Bellman, S., & Lohse, G. L. (2002). Defaults, framing and privacy: Why opting in-opting out. *Marketing Letters*, *13*(1), 5–15. https://doi.org/10.1023/A:1015044207315
15. Levav, J., Heitmann, M., Herrmann, A., & Iyengar, S. (2010). Order in product customization decisions: Evidence from field experiments. *Journal of Political Economy*, *118*(2), 274–299. https://doi.org/10.1086/652463
16. Brown, Z., Johnstone, N., Hascic, I., Vong, L., & Barascud, F. (2013). Testing the effect of defaults on the thermostat settings of OECD employees. *Energy Economics*, *39*, 128–134. https://doi.org/10.1016/j.eneco.2013.04.011
17. Haggag, K., & Paci, G. (2014). Default tips. *American Economic Journal: Applied Economics*, *6*(3), 1–19. https://doi.org/10.1257/app.6.3.1
18. Patel, M. S., Day, S., Small, D. S., Howell, J. T., III, Lautenbach, G. L., Nierman, E. H., & Volpp, K. G. (2014). Using default options within the electronic health record to increase the prescribing of generic-equivalent medications. *Annals of Internal Medicine*, *161*(10, Suppl.), S44–S52. https://doi.org/10.7326/M13-3001
19. Patel, M. S., Day, S. C., Halpern, S. D., Hanson, C. W., Martinez, J. R., Honeywell, S., Jr., & Volpp, K. G. (2016). Generic medication prescription rates after health system-wide redesign of default options within the electronic health record. *JAMA Internal Medicine*, *176*(6), 847–848. https://doi.org/10.1001/jamainternmed.2016.1691
20. Goswami, I., & Urminsky, O. (2016). When should the ask be a nudge? The effect of default amounts on charitable donations. *Journal of Marketing Research*, *53*(5), 829–846. https://doi.org/10.1509/jmr.15.0001
21. Altmann, S., Falk, A., Heidhues, P., Jayaraman, R., & Teirlinck, M. (2019). Defaults and donations: Evidence from a field experiment. *The Review of Economics and Statistics*, *101*(5), 808–826. https://doi.org/10.1162/rest_a_00774
22. Jachimowicz, J. M., Duncan, S., Weber, E. U., & Johnson, E. J. (2019). When and why defaults influence decisions: A meta-analysis of default effects. *Behavioural Public Policy*, *3*(2), 159–186. https://doi.org/10.1017/bpp.2018.43
23. Carroll, G. D., Choi, J. J., Laibson, D., Madrian B. C., & Metrick, A. (2009). Optimal defaults and active decisions. *The Quarterly Journal of Economics*, *124*(4), 1639–1674. https://doi.org/10.1162/qjec.2009.124.4.1639
24. Keller, P. A., Harlam, B., Loewenstein, G., & Volpp, K. G. (2011). Enhanced active choice: A new method to motivate behavior change. *Journal of Consumer Psychology*, *21*(4), 376–383. https://doi.org/10.1016/j.jcps.2011.06.003
25. Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (2021). Active choice, implicit defaults, and the incentive to choose. *Organizational Behavior and Human Decision Processes*, *163*, 6–16. https://doi.org/10.1016/j.obhdp.2019.02.001
26. Bernheim, B. D., Fradkin, A., & Popov, I. (2015). The welfare economics of default options in 401(k) plans. *American Economic Review*, *105*(9), 2798–2837. https://doi.org/10.1257/aer.20130907
27. McKenzie, C. R. M., Liersch, M. J., & Finkelstein, S. R. (2006). Recommendations implicit in policy defaults. *Psychological Science*, *17*(5), 414–420. https://doi.org/10.1111/j.1467-9280.2006.01721.x
28. Dinner, I., Johnson, E. J., Goldstein, D. G., & Liu, K. (2011). Partitioning default effects: Why people choose not to choose. *Journal of Experimental Psychology: Applied*, *17*(4), 332–341. https://doi.org/10.1037/a0024354
29. Fowlie, M., Wolfram, C., Spurlock, C. A., Todd-Blick, A., Baylis, P., & Cappers, P. (2020). *Default effects and follow-on behavior: Evidence from an electricity pricing program* (Working Paper 23553). National Bureau of Economic Research. https://doi.org/10.3386/w23553
30. Pope, D., & Simonsohn, U. (2011). Round numbers as goals: Evidence from baseball, SAT takers, and the lab. *Psychological Science*, *22*(1), 71–79. https://doi.org/10.1177/0956797610391098