How to Choose a Default

How to Choose a Default

John Beshears (corresponding author)
Harvard Business School and National Bureau of Economic Research
Soldiers Field, Boston, MA 02163
jbeshears@hbs.edu

Richard T. Mason
Voya Financial and Carnegie Mellon University
One Orange Way, Windsor, CT 06095
richardt.mason@voya.com

Shlomo Benartzi
University of California, Los Angeles
110 Westwood Plaza, Los Angeles, CA 90095
benartzi@savemoretomorrow.org

March 25, 2021

Abstract: Consider the default that applies to a group of individuals in a setting featuring many ordered options. We present a framework for understanding how the default influences the group's distribution of outcomes. A key input for the framework is the distribution of latent options—the options that individuals would choose if they were forced to make a choice in the absence of a default. In the framework, the default acts as an attraction point. For individuals whose latent options are in a neighborhood of the default, the default pulls their outcomes to be closer to the default. To provide support for the assumptions of the framework, we report the results of new field experiments in the domain of retirement savings, and we discuss the relationship between the framework and prior empirical evidence on defaults. The framework guides policy makers who have the opportunity to select a default.

Keywords: nudge, choice architecture, behavioral economics, behavioral science, default, savings

---

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, Giovanni Urga, and four anonymous editors and reviewers. Beshears is an associate professor at Harvard University and a faculty research fellow at NBER. Mason is an employee of Voya Financial and a visiting scholar and Senior Research Fellow at Carnegie Mellon University. Benartzi is a professor emeritus at University of California, Los Angeles and 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 a complete list of outside activities. Voya Financial had the opportunity to review the manuscript before public release for the sake of identifying factual inaccuracies, but the authors retained full editorial control of the manuscript contents. The findings and conclusions expressed are solely those of the authors and do not represent the views of Voya Financial, Harvard University, NBER, Carnegie Mellon University, or University of California, Los Angeles.

1. Introduction

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, offered two retirement savings policies, and the first was "changing the default" from non-participation to participation in defined contribution plans.[1,A] Past research has documented that moving from an opt-in plan—which implements a default savings rate of zero for individuals who do not actively elect a different rate—to an opt-out plan—which features a strictly positive default savings rate—increases plan participation dramatically.[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 outcomes 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 documents the effects of defaults across settings.[22]

In this paper, we offer guidance to policy makers who must choose a default among many ordered options.[B] We present a framework for understanding the effects of defaults. We label the framework an "empirical model" because it does not specify the economic and psychological mechanisms driving individuals' responses to defaults. It merely describes how different defaults shape the distribution of outcomes. To match our experimental evidence and to use concrete terminology, we discuss the model using the language of defined contribution retirement savings plans, such as 401(k) plans.

A key input to the model is the distribution of "latent" contribution rates, which we define as the contribution rates that individuals would select if they were forced to choose in the absence of a default. In the model, the default is an attraction point. 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 contribution 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 person's outcome.

Data from new field experiments, which we conducted prior to developing the model, provide empirical evidence consistent with the model. An important feature of our primary experiment (Study 1) is that one condition asks individuals to select contribution rates in the absence of a default, revealing the distribution of latent contribution rates.[20-21,23-25] Study 1 also features a *7% default* condition and a *10% default* condition. The default conditions do not increase the mean contribution rate relative to the *no default* condition, but this does not imply that the defaults have no influence on outcomes. The defaults change the distribution of contribution rates in ways that are consistent with the model, for example by increasing the fraction of individuals with contribution rates equal to the default. Our findings are also consistent with previous research, and we discuss how the empirical evidence on defaults aligns with the model assumptions. We calibrate the parameters of the model using our experimental data and find that the influence of defaults is weaker in our setting than in other settings.

To draw out the implications of the model, we take the perspective of a policy maker who generally wishes to shift contribution rates upwards.[C] The policy maker may believe, for example, that a group of individuals is saving too little and should increase their retirement plan contributions.[8] This perspective distinguishes our paper from prior work on how to set defaults,[23,26] and analogous perspectives may apply in other settings. Perhaps a group of workers sets their office thermostats to too high a temperature during the winter and should generally decrease their thermostat settings,[16] or perhaps a group of physicians is overprescribing brand-name medications and should generally switch to prescribing generic equivalents.[18-19]

According to our model, the policy maker must balance several considerations. When the default is above an individual's latent contribution rate, increasing the default will pull the individual's contribution rate higher, conditional on influencing that individual. However, holding fixed the individual's latent contribution rate, a higher default is less likely to influence that individual. At the same time, the default may pull down the contribution rates of individuals whose latent contribution rates are higher than the default.

The policy maker who is seeking to increase the mean contribution rate should select a default that pulls many contribution rates up while pulling few down. In general, the policy maker 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 only slightly above that cluster. When the default has a strong influence, the default that maximizes the mean contribution rate is likely to be further above the cluster.

While this guidance presumes that the policy maker has knowledge of both the distribution of latent contribution rates and the strength of the default, the final section of the paper also addresses cases in which the policy maker's ability to gather this information is limited. In addition, the final section discusses caveats to our analysis and factors outside the model that the policy maker should consider.

2. The Model

Our empirical model has three ingredients: (1) the distribution of latent contribution rates, (2) the location of the default, and (3) a formula governing whether and to what extent the default affects an individual with a given latent contribution rate.

To make the terminology precise, consider our experimental setting. An individual visits a website for selecting a contribution rate. When there is no default, the individual uses the keyboard to enter a number into a blank space. Then, the individual can either proceed with that initial number or take steps to select a different number, whether by interacting with the website further or by returning to the website soon thereafter. An individual's latent contribution rate is defined as the individual's contribution rate at the end of this process.

When a default is in place, the individual visits the website and sees that the space has already been filled in with the default contribution rate. Again, the individual can either proceed with that initial number or take steps to select a different number.

The heart of the model is the third ingredient, which is decomposed into two steps. Given the default and an individual's latent contribution rate, what is the likelihood that the individual is affected by the default? Then, conditional on being affected, what is the individual's ultimate contribution rate?

Given an individual's latent contribution rate L and the default D, we assume that the probability that the individual is influenced by the default is generally decreasing in the absolute distance between L and D, but that there are certain latent contribution rates that individuals find particularly attractive. We see in our experiments and in other data sets that individuals like to choose contribution rates that are multiples of five, so we define the function $A(L)$, the adjustment factor for attractive contribution rates, to be equal to the parameter F (where $0 \leq F \leq 1$) when L is a multiple of five and equal to zero otherwise. Then, we define the parameter R to be the radius within which the default has an effect. An individual whose latent contribution rate is equal to the default is 100% likely to be influenced by the default. An individual whose latent contribution rate is a distance of exactly R away from the default has 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 a multiple of five. In mathematical notation, for $|D - L| \leq R$, the probability that the individual is influenced by the default is $(1 - A(L)) \times (1 - |D - L| / R)$. The probability is zero for $|D - L| > R$. Note that we set the probability for $|D - L| > R$ equal to zero for the sake of simplifying the model. In reality, even individuals whose latent contribution rates are very far from the default are likely to have a non-zero probability of being influenced by the default. Our simplifying assumption is meant to capture the notion that this probability is likely to be much closer to zero than the corresponding probability for an individual whose latent contribution rate is close to the default.

We assume that conditional on being affected by the default, an individual ends up with the contribution rate $((1 - W)L + WD)$, where $0 \leq W \leq 1$. This is a weighted average of the latent contribution rate and the default. Because the online interface in our experiment encouraged individuals to choose contribution rates as whole number percentages, we round the contribution rate given by the formula to the nearest whole number.

To illustrate how the model works, we construct two examples. Figure 1 shows the first example. The white bars display a distribution of latent contribution rates, which is the distribution of contribution rates observed in our primary experiment when there is no default. The grey and black bars show the distribution of contribution rates when the default is 7% or 10%, respectively, as predicted by the model with R (the radius within which the default has an effect) taking a value of 12, F (the value of the adjustment factor for contribution rates that are multiples of five) taking a value of 0.3, and W (the weight placed on the default among individuals affected by the default) taking a value of 0.9. These values imply that the default has a strong impact, and the grey and black bars show that a wide swathe of the distribution is drawn towards the default.

Figure 2 shows the second example, which is the same as the first except that we use R = 1.5, F = 0.3, and W = 0.7. As we explain later in the paper, these values give the best fit for our data. They imply a weaker default effect, but the grey bars indicate that the 7% default still draws in individuals whose latent contribution rates are 6% or 8%. The black bars show that the 10% default has less of an impact on the distribution. It draws in individuals whose latent contribution rates are 9% or 11%, but few individuals have those latent contribution rates.

As mentioned in the Introduction, we construct our empirical model without specifying the mechanisms driving individuals' responses to defaults. Nonetheless, our model can be compared to existing models that do articulate mechanisms for default effects. Our model implies that individuals whose latent contribution rates are closer to the default are more likely to be impacted by the default, and this feature is shared with models that assume individuals face a cost of opting out of the default.[21,23,26] In these models, individuals' latent contribution rates are assumed to be their most preferred contribution rates, and individuals whose latent contribution rates are far from the default therefore have a stronger incentive to incur the opt-out cost and switch to their most preferred contribution rates. Individuals whose latent contribution rates are close to the default have a weaker incentive to incur the opt-out cost and are more likely to remain at the default. At the same time, our model is distinct from models involving opt-out costs because our model allows for the possibility that the default increases the probability that individuals choose contribution rates close to but not equal to the default. In a model of opt-out costs, individuals who incur the opt-out cost choose their latent contribution rates, so the default does not increase the prevalence of contribution rates close to but not equal to the default. An anchoring model, on the other hand, does allow for the default to cause such an increase.[26] As we will describe in the next section, however, we do not observe such an increase in our empirical setting, so our data cannot distinguish among mechanisms underlying default effects.

3. Evidence

3.A. Experimental Design

We conducted three field experiments in collaboration with Voya, a U.S. retirement services and recordkeeping provider. We focus on Study 1 because it includes a condition with no default, allowing us to observe the distribution of latent contribution rates. In this subsection, we describe Study 1. We describe Studies 2 and 3, which were similar, in a later subsection. For details, see the Supplemental Material.

We worked with the segment of Voya 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. SM Figures 1-8 show screenshots.

Randomization occurred when employees reached the webpage on which they selected their contribution rates. Study 1 featured three conditions. In the *7% default* and *10% default* conditions, the space for indicating the desired contribution rate was prepopulated with a default of 7% or 10%, respectively (SM Figure 4). In the *no default* condition, the space for indicating the desired contribution rate was empty when the webpage loaded, and a blinking cursor suggested that the employee should enter a number (SM Figure 5). 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 SM Figure 4). In all three conditions, employees could increase or decrease their chosen contribution rate away from the initial rate by clicking on "+" or "-" buttons.

As specified when we pre-registered Study 1 (SM Figure 9), our primary outcome variable is the contribution rate in effect 60 days after the initial Voya Enroll visit, winsorized (to reduce the 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. The choice of a 60-day window balances two factors. First, using too long a window introduces substantial noise into our measurement of default effects. A range of factors, such as salary increases, financial emergencies, and changes in the employment status of family members, cause individuals to change their contribution rates over time and make it more difficult to observe the effect of defaults at long horizons. On the other hand, using too short a window, for example by analyzing contribution rates one day after the initial Voya Enroll visit, fails to capture adjustments that employees make as they respond to defaults. Some employees choose not to enroll in the plan when they first visit Voya Enroll, but they return a few days later and select a positive contribution rate.

3.B. Experimental Results and Prior Literature Supporting the Model Assumptions

As shown in Table 1, approximately half of the individuals in Study 1 were male. The mean age was 38, and the mean annual salary was approximately $70,000. These characteristics were not statistically significantly different across the three conditions.

Following our pre-registered analysis plan, we perform ordinary least squares regressions with contribution rate as the outcome variable and indicators for experimental condition as the explanatory variables of interest. The point estimates indicate that relative to having no default, the 7% default decreased the mean contribution rate by 0.02 percentage points when omitting controls for gender, age, and salary and by 0.04 percentage points when including controls, while the 10% default increased the mean contribution rate by 0.08 percentage points when omitting controls and by 0.06 percentage points when including controls. None of these estimates are statistically significant, and all of them are small in magnitude. When we use the same regression framework to investigate whether the 7% default and 10% default increased the likelihood that an individual has a contribution rate greater than zero, we similarly find estimates that are small in magnitude and not statistically significant.

We had hypothesized that the 7% default and 10% default would increase the mean contribution rate relative to having no default, so we were surprised by these results. However, the lack of an effect on the mean contribution rate does not imply that the defaults did not influence contribution rates. The analyses reported below, which were not pre-registered, indicate that the defaults indeed had impacts on the distribution of contribution rates. We use these findings to offer empirical support for our model assumptions.

Figure 3 shows the contribution rate distributions in the experimental conditions. It is immediately apparent that individuals are attracted to contribution rates that are multiples of five. This finding accords with previous work[4] and is the reason for the assumption in our model that individuals whose latent contribution rates are multiples of five may be less likely to be influenced by defaults.[D]

Figure 4 formally compares the contribution rate distributions in the *7% default* and *10% default* conditions to the *no default* condition. We conduct a series of ordinary least squares regressions where the outcome variable is an indicator for the contribution rate being equal to C%, with C taking on each integer value from zero to 15. We also conduct a regression where the outcome variable is an indicator for the contribution rate being greater than or equal to 16%. The explanatory variables are indicators for the *7% default* and *10% default* conditions. In Figure 4, the value for C varies along the horizontal axis. The vertical bars indicate the point estimates for the effect of the 7% default or the 10% default on the likelihood of having a contribution rate equal to C%, relative to having no default. The whiskers give 95% confidence intervals.

Figure 4 indicates that relative to having no default, the 7% default causes a statistically significant increase in the fraction of individuals with a 7% contribution rate, and the 10% default causes a statistically significant increase in the fraction of individuals with a 10% contribution rate. This finding is consistent with our model, which implies that some individuals with latent contribution rates close to the default end up choosing the default. This finding is also consistent with the prior literature on default effects, which documents that the default is chosen frequently.[2-6,17,20-21]

The model further implies that individuals sometimes end up choosing the default both when their latent contribution rates are below the default and when their latent contribution rates are above the default. Consistent with this prediction, the 7% default decreases the fraction of employees with a contribution rate less than or equal to 6% and decreases the fraction of employees with a contribution rate greater than or equal to 8%. The 10% default decreases the fraction of employees with a contribution rate less than or equal to 9% and decreases the fraction of employees with a contribution rate greater than or equal to 11%, although the last estimate is not statistically significant, perhaps due to a floor effect. These results are consistent with previous findings.[5-6,20-21]

Finally, the model assumes that the 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. This assumption is consistent with past research.[5-6,17,20-21] Additional analyses of the Study 1 data also provide suggestive evidence that is consistent with the assumption, as we describe in more detail in Section E of the Supplemental Material, although it is important to emphasize that the statistical power of these tests is low.[E] As we discuss in the next subsection, the calibration of our model similarly suggests that defaults indeed have an influence only on individuals whose latent contribution rates are close to the default.

3.C. Model Calibration

To calibrate our model, we combine data from Study 1 with data from Studies 2 and 3. Studies 2 and 3 were conducted prior to Study 1 and were not pre-registered. They have the same design as Study 1 except that they do not have a *no default* condition and instead have conditions with integer contribution rate defaults 6% through 11%. See the Supplemental Material.

Given particular values of the model parameters R, F, and W, and given the distribution of latent contribution rates from the *no default* condition of Study 1, we calculate the model's predictions for the distribution of contribution rates with a 6%, 7%, 8%, 9%, 10%, or 11% default. Then, we compare the model's predictions to the observed distributions of contribution rates in Studies 1, 2, and 3, and we calculate a summary measure of the extent to which the predicted and observed distributions differ. The model best fits the data when the parameters take a combination of values that minimizes the summary measure of differences, as described in detail in the Supplemental Material.

The model best fits the data when R (the radius within which the default has an effect) takes a value of 1.5, F (the value of the adjustment factor for contribution rates that are multiples of five) takes a value of 0.3, and W (the weight placed on the default among individuals affected by the default) takes a 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 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, 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 choose contribution rates of 9% or more, relative to the *no default* condition. Similarly, the *10% default* condition leads to a statistically significant 2.2 percentage point decrease in the fraction of individuals who choose contribution rates of 8% or less, relative to the *no default* condition. (The *10% default* condition does not have a statistically significant effect on the fraction of individuals who choose contribution rates of 12% or more, relative to the *no default* condition). With these caveats, the best-fitting parameter values for the model include a low value of R, indicating that the default tends to attract individuals whose latent contribution rates are close to the default.

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.[F]

4. 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, the discussion in subsection 3.B indicates that 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 the policy maker is trying to increase the mean outcome, our model provides guidance for selecting a default. First, the policy maker should consider the distribution of latent outcomes. Next, the policy maker should gauge how influential the default is. With this information, the policy maker should place the default such that it pulls up the outcomes of many individuals while pulling down the outcomes of few individuals. If the default is weak (R, the radius 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 (R is large), the default will likely be higher. When F (the value of the adjustment factor for focal options in the choice set) is high, the policy maker must be wary of latent outcomes that individuals are reluctant to leave.

For a policy maker 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, non-experimental data are informative. If few individuals end up with the default option, the influence of the default is weak, and the distribution of observed outcomes gives an approximation to the distribution of latent outcomes. If many individuals end up with the default option, the influence of the default is strong. In this case, it is difficult to learn the distribution of latent outcomes, but the prescription is that the policy maker should push the default to be more extreme in order to shift outcomes in the desired direction.

For additional insight into the likely strength of the default, a policy maker can rely on past research. According to prior work, defaults are more impactful in consumer domains and less impactful in environmental domains, and they are more influential when they communicate the policy maker's recommendation[27] or serve as a reference point against which other options are judged[28] than when they make the default option easy to implement.[22]

It is important to note limitations of our analysis. The model applies to our experimental setting and to many other settings, but it does not apply to all situations. 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 assumptions of the model may not be satisfied.[29] In these situations, it is less likely true that the influence of the default is weaker as the difference between the default and the individual's latent outcome increases. This observation highlights a key feature of our experimental setting. The individuals in our experiment made a choice to visit the 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 observe are weaker than some other default effects that have been documented previously in retirement savings plans.[2-6] Perhaps the individuals in our experiment arrived at the website having already contemplated the contribution rate they would like to choose, diminishing the scope for the default to influence this decision.

Our paper has not addressed the moral considerations that a policy maker should have in mind when choosing a default. We have adopted the perspective of a policy maker who is trying to shift outcomes in one direction under the assumption that the policy maker has ethically sound reasons for doing so. For example, the policy maker may have strong reasons to believe that psychological biases are causing individuals' choices to deviate systematically from the choices that would maximize their own welfare. For another example, the policy maker may wish to shift outcomes because individuals are making decisions in ways that do not account for externalities imposed on others. Policy makers who are less sure which outcomes are appropriate should use a different framework for contemplating default selection.[26] They should be wary of 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 that have different model parameters and different latent outcome distributions, it would be possible to assign each subpopulation to its own tailored default. This line of reasoning can be applied to situations in which the policy maker has a more complex objective than simply shifting outcomes upwards or downwards. For example, if a policy maker 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 observed income, with one group's default chosen to increase contribution rates and the other group's default chosen to promote more moderate contribution rates. If the policy maker's objective depends on an individual's unobserved attributes, it is interesting to ask whether the policy maker could elicit information about those attributes before assigning a default.

As another extension, it would be valuable to consider how a default might change over time. For example, consider the case of a smart thermostat. To reduce energy consumption, the thermostat might initially have a default temperature that is only slightly below the temperature that individuals would choose for themselves during winter. After a period of time, as individuals habituate to colder temperatures, the thermostat might have a lower default temperature.

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

Endnotes

A. The second policy Thaler mentioned was automatic escalation of contribution rates in retirement plans, which is similar to the first because it sets future defaults.

B. We do not address situations featuring a small number of options (say, five or fewer) in the 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, as they do not map onto the structure of our model.

C. The implications of the model are symmetric for a policy maker who generally wishes to shift contribution rates downwards.

D. One could make the argument that these people 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 subsection 3.C, our calibration exercise indicates that giving F a value of 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 Pope, D., & Simonsohn, U. (2011). Round numbers as goals: Evidence from baseball, SAT takers, and the lab. *Psychological Science, 22,* 71-79.

E. The functional form assumptions in the model impose symmetry between the effect of the default on individuals with lower latent contribution rates and the effect on those with higher latent contribution rates. The additional analyses described in Section E of the Supplemental Material provide a test of this assumption. 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.

F. Note, however, that when we pool the data from Studies 2 and 3, we find that the 7% default leads to a higher mean contribution rate than the 6% default, an effect that is marginally statistically significant (SM Table 4). Given the confidence interval around this estimate and the number of experimental conditions, we do not view this result as contradictory to our model.

References

1. Dubner, S. J. (2018), People aren't dumb. The world is hard. *Freakonomics Radio, episode 340,* https://freakonomics.com/podcast/thaler-rebroadcast, accessed August 17, 2020.

2. Madrian, B. C., & Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics, 116,* 1149-1187.

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, 16* (pp. 67-113). Chicago, IL: University of Chicago Press.

4. Choi, J. J., Laibson, D., & Madrian, B. C. (2004). Plan design and 401(k) savings outcomes. *National Tax Journal, 57,* 275-298.

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). Chicago, IL: University of Chicago Press.

6. Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (2008). The importance of default options for retirement saving outcomes: Evidence from the United States. In S. J. Kay and T. Sinha (Eds.), *Lessons from pension reform in the Americas* (pp. 59-87). Oxford: Oxford University Press.

7. Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness.* New Haven, CT: Yale University Press.

8. Benartzi, S., & Thaler, R. H. (2013). Behavioral economics and the retirement savings crisis. *Science, 339,* 1152-1153.

9. Plan Sponsor Council of America (2018). *60th Annual Survey of Profit Sharing and 401(k) Plans.* Chicago, IL: Plan Sponsor Council of America.

10. Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science, 302,* 1338-1339.

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,* 599-620.

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,* 69-88.

13. Johnson, E. J., Hershey, J., Meszaros, J., & Kunreuther, H. (1993). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty, 7,* 35-51.

14. Johnson, E. J., Bellman, S., & Lohse, G. L. (2002). Defaults, framing and privacy: Why opting in-opting out. *Marketing Letters, 13,* 5-15.

15. Levav, J., Heitmann, M., Herrmann, A., & Iyengar, S. (2010). Order in product customization decisions: Evidence from field experiments. *Journal of Political Economy, 118,* 274-299.

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.

17. Haggag, K., & Paci, G. (2014). Default tips. *American Economic Journal: Applied Economics, 6(3),* 1-19.

18. Patel, M. S., Day, S., Small, D. S., Howell III, J. T., 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,* S44-S52.

19. Patel, M. S., Day, S. C., Halpern, S. D., Hanson, C. W., Martinez, J. R., Honeywell Jr., S., & 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,* 847-848.

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,* 829-846.

21. Altmann, S., Falk, A., Heidhues, P., Jayaraman, R., & Teirlinck, M. (2019). Defaults and donations: Evidence from a field experiment. *Review of Economics and Statistics, 101,* 808-826.

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,* 159-186.

23. Carroll, G. D., Choi, J. J., Laibson, D., Madrian B. C., & Metrick, A. (2009). Optimal defaults and active decisions. *Quarterly Journal of Economics, 124,* 1639-1674.

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,* 376-383.

25. Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (forthcoming). Active choice, implicit defaults, and the incentive to choose. *Organizational Behavior and Human Decision Processes.*

26. Bernheim, B. D., Fradkin, A., & Popov, I. (2015). The welfare economics of default options in 401(k) plans. *American Economic Review, 105,* 2798-2837.

27. McKenzie, C. R. M., Liersch, M. J., & Finkelstein, S. R. (2006). Recommendations implicit in policy defaults. *Psychological Science, 17,* 414-420.

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,* 332-341.

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.

Table 1. Gender, Age, and Salary by Randomly Assigned Condition: Study 1

	Experimental condition			p-value from chi-squared or F test for null hypothesis that conditions are equal
	No default	7% default	10% default	
Percentage male	53	52	52	0.66
Mean age (standard deviation)	38 (12)	38 (12)	38 (12)	0.69
Mean salary ($000s) (standard deviation)	69 (51)	71 (52)	71 (54)	0.16
Observations	3,991	4,024	4,048	

Figure 1. Illustration of the Model with a Strong Default (*R* = 12, *F* = 0.3, *W* = 0.9)

The white bars show the distribution of contribution rates in the *no default* condition in Study 1. The grey and black bars show the distribution of contribution rates predicted by the model for a 7% default and for a 10% default, respectively, using the distribution of contribution rates in the *no default* condition in Study 1 as the distribution of latent contribution rates and using the parameter values R = 12, F = 0.3, and W = 0.9.

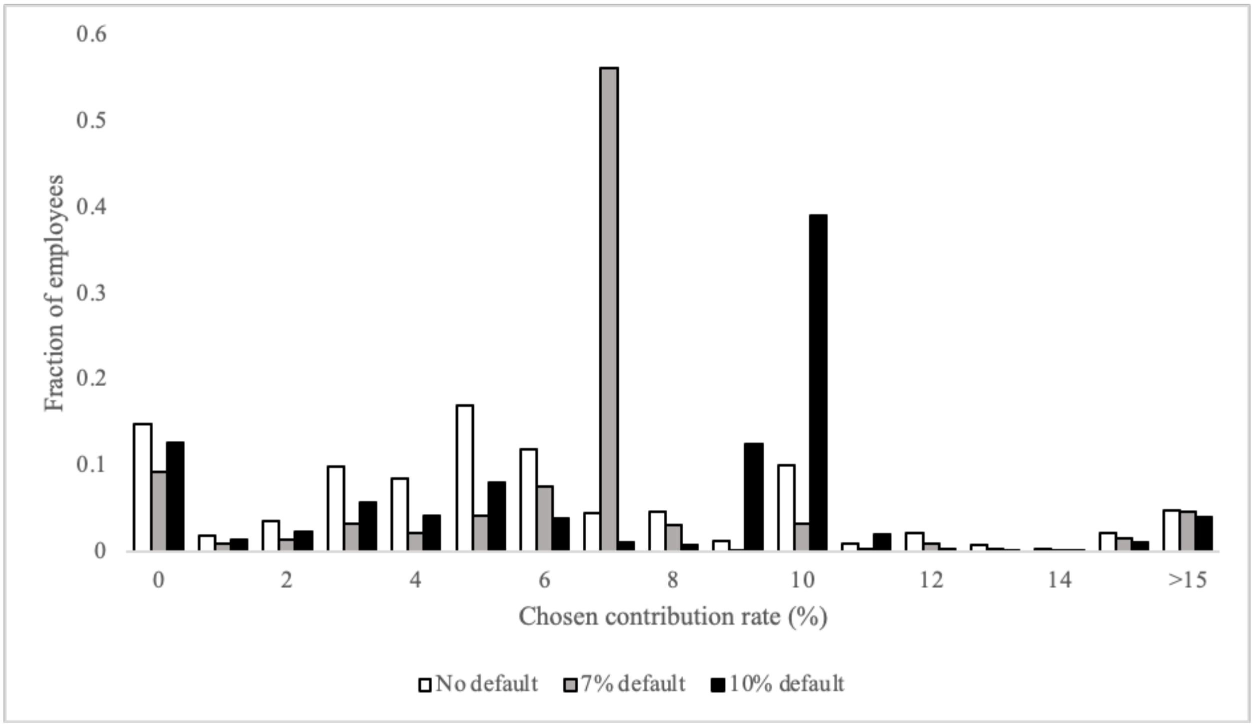

Figure 2. Illustration of the Model with a Weak Default (*R* = 1.5, *F* = 0.3, *W* = 0.7)

The white bars show the distribution of contribution rates in the *no default* condition in Study 1. The grey and black bars show the distribution of contribution rates predicted by the model for a 7% default and for a 10% default, respectively, using the distribution of contribution rates in the *no default* condition in Study 1 as the distribution of latent contribution rates and using the parameter values $R = 1.5$, $F = 0.3$, and $W = 0.7$. These parameter values are the best fit for the experimental data (see subsection 3.C).

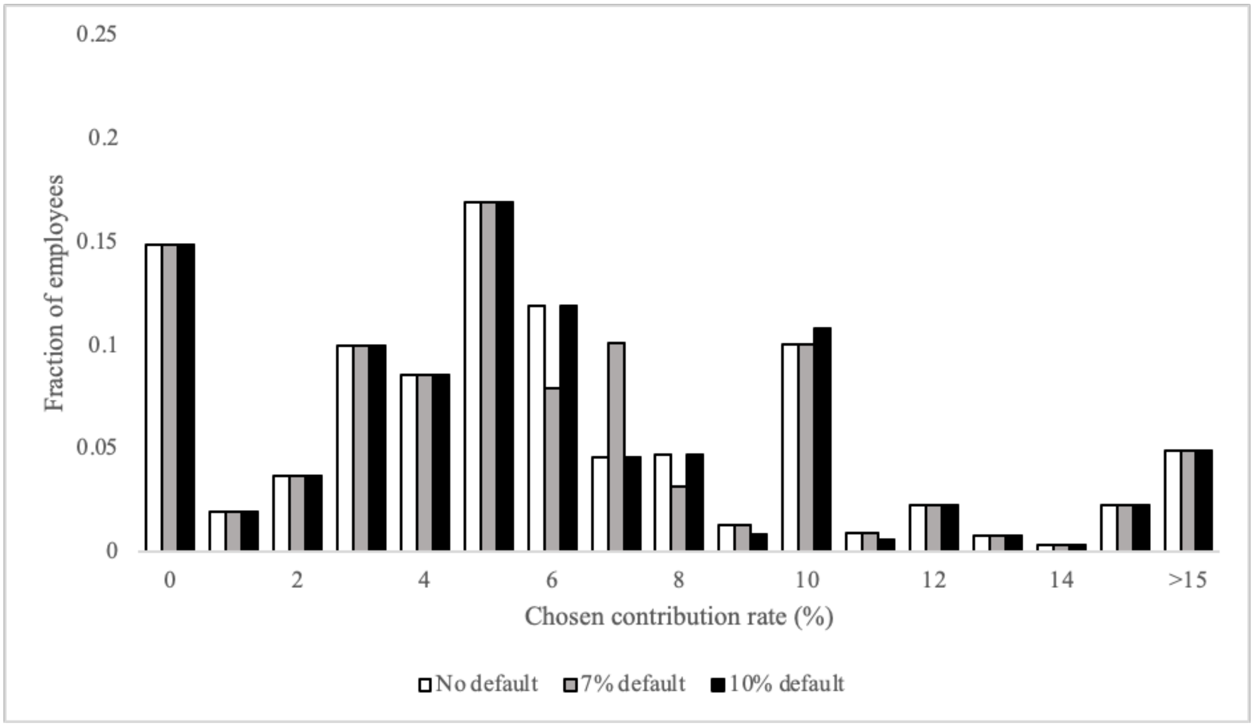

Figure 3. Distributions of Contribution Rates by Randomly Assigned Condition: Study 1

The white, grey, and black bars show the distribution of contribution rates in Study 1 in the *no default* condition, in the *7% default* condition, and in the *10% default* condition, respectively. Contribution rates are measured 60 days after the initial visit to the retirement plan enrollment website.

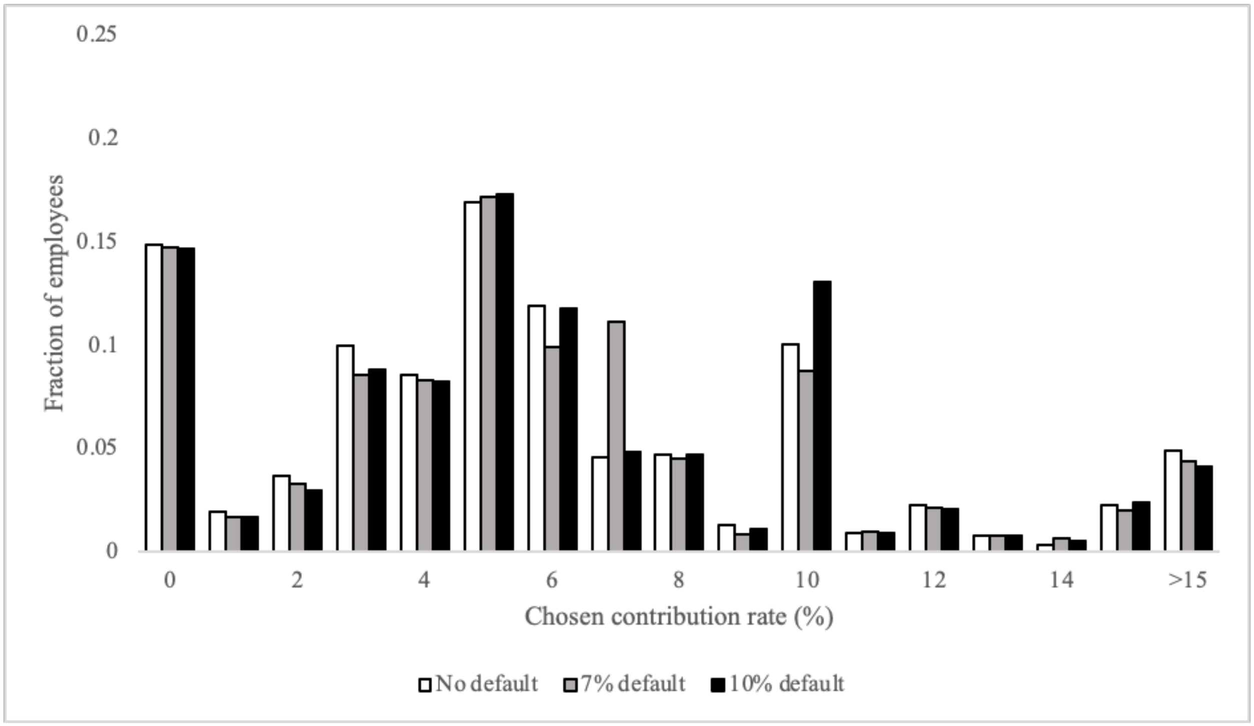

Figure 4. The Effect of the 7% Default and the 10% Default on the Likelihood of the Contribution Rate Being Equal to a Given Value: Study 1

We conduct a series of ordinary least squares regressions where the outcome variable is an indicator for the contribution rate being equal to *C*%, with *C* taking on each integer value from zero to 15. We also conduct a regression where the outcome variable is an indicator for the contribution rate being greater than or equal to 16%. The explanatory variables are indicators for the *7% default* and *10% default* conditions. The value for *C* varies along the horizontal axis. The vertical bars indicate the point estimates for the effect of the 7% default or the 10% default on the likelihood of having a contribution rate equal to *C*%, relative to having no default. The whiskers give 95% confidence intervals.

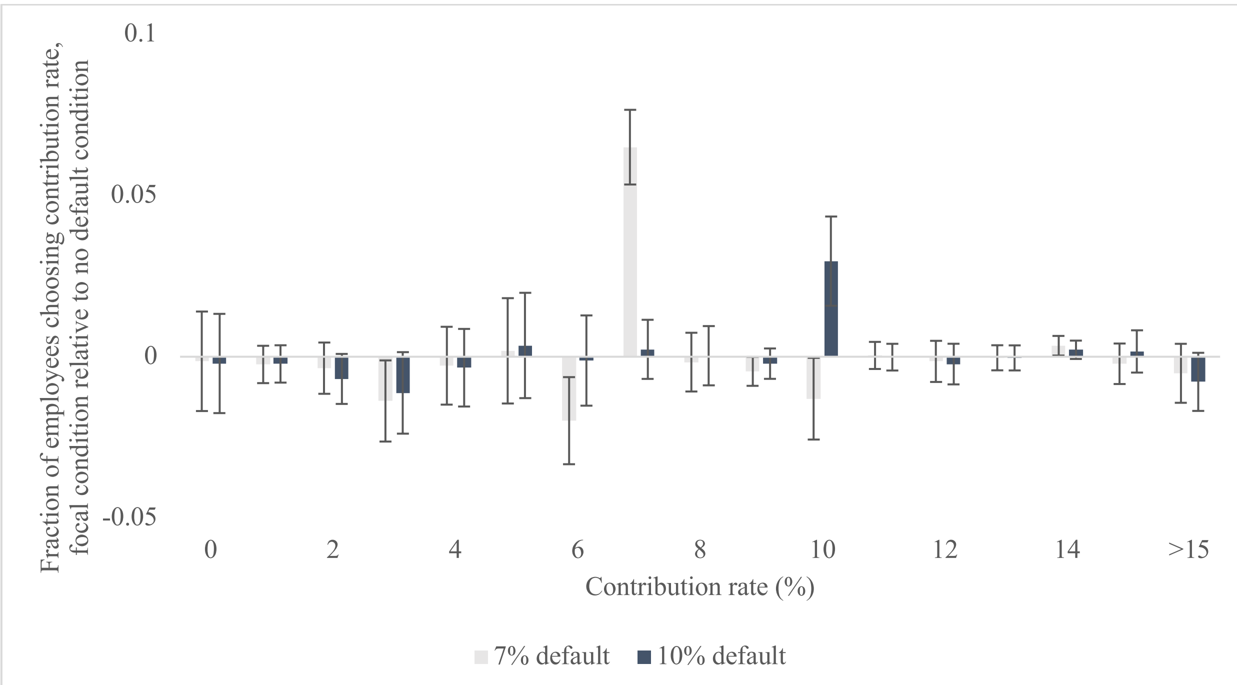

Figure 5. Mean Contribution Rate Predicted by the Model with Best-Fit Parameter Values (*R* = 1.5, *F* = 0.3, *W* = 0.7) as the Default Varies

For each integer default contribution rate from zero to 15, we calculate the model's prediction for the distribution of individual contribution rates and the mean of that distribution. As inputs to the model, we use the distribution of contribution rates in the *no default* condition in Study 1 as the distribution of latent contribution rates, and we use the parameter values $R = 1.5$, $F = 0.3$, and $W = 0.7$. These parameter values are the best fit for the experimental data (see subsection 3.C). The default contribution rate varies along the horizontal axis, and the vertical axis is the model-predicted mean contribution rate. The horizontal line in the middle of the figure shows the mean contribution rate in the *no default* condition in Study 1.