

ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

'BETTER OFF, AS JUDGED BY THEMSELVES': BOUNDED RATIONALITY AND NUDGING

Cass R. Sunstein

Forthcoming in *Routledge Handbook on Bounded Rationality*

Discussion Paper No. 942

11/2017

Harvard Law School
Cambridge, MA 02138

This paper can be downloaded without charge from:

The Harvard John M. Olin Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center

The Social Science Research Network Electronic Paper Collection:
<https://ssrn.com/abstract=3059310>

“Better Off, As Judged By Themselves”: Bounded Rationality and Nudging

Cass R. Sunstein*

Abstract

Many nudges are designed to make people better off, as judged by themselves. This criterion, meant to ensure that nudges will increase people’s welfare, contains some ambiguity. It is useful to distinguish among three categories of cases: (1) those in which choosers have clear antecedent preferences, and nudges help them to satisfy those preferences (often by increasing “navigability”); (2) those in which choosers face a self-control problem, and nudges help them to overcome that problem; and (3) those in which choosers would be content with the outcomes produced by two or more nudges, or in which ex post preferences are endogenous to nudges, so that without additional clarification or work, the “as judged by themselves” criterion does not identify a unique solution for choice architects. Category (1) is self-evidently large. Because many people agree that they suffer self-control problems, category (2) is large as well. Cases that fall in category (3) create special challenges, which may lead us to make direct inquiries into welfare or to explore what informed, active choosers typically select.

Some nudges are designed to reduce externalities; consider fuel economy labels that draw attention to environmental consequences, or default rules that automatically enroll people in green energy (Reisch and Sunstein 2014). But many nudges are designed to increase the likelihood that people’s choices will improve their own welfare. Richard Thaler and I argue that the central goal of such nudges is to “make choosers better off, *as judged by themselves*” (Thaler and Sunstein, 2008, p. 5; italics in original; Thaler, 2015). Social planners – or in our terminology, choice architects – might well have their own ideas about what would make choosers better off, but in our view, the lodestar is people’s own judgments. To be a bit more specific: The lodestar is welfare, and people’s own judgments are a good (if imperfect) way to test the question whether nudges are increasing their welfare.

The last sentence raises many questions, and it is certainly reasonable to wonder about potential ambiguities in the “as judged by themselves” (hereinafter AJBT) criterion. Should the focus be on choosers’ judgments before the nudge, or instead after? What if the nudge alters people’s preferences, so that they like the outcome produced by the nudge, when they would not have sought that outcome in advance? What if preferences are constructed by the relevant choice

* Robert Walmsley University Professor, Harvard University. I am grateful to the Program on Behavioral Economics and Public Policy for support and to Jacob Goldin and Lucia Reisch for valuable comments on a previous draft. An earlier version of this essay appeared in the *International Review of Economics* (2017), available at <https://link.springer.com/article/10.1007%2Fs12232-017-0280-9>

architecture (Lichtenstein and Slovic 2006)? Or what if people's ex ante judgments are wrong, in the sense that a nudge would improve their welfare, even though they do not think that it will (Dolan 2014)? Do we want to ask about choosers' actual, potentially uninformed or behaviorally biased judgments, or are we entitled to ask what choosers would think if they had all relevant information and were unaffected by relevant biases (Goldin 2015)?

My goal here is to explore the meaning of the AJBT criterion and to sort out some of the ambiguities. As we shall see, three categories of cases should be distinguished: (1) those in which choosers have clear antecedent preferences, and nudges help them to satisfy those preferences; (2) those in which choosers face a self-control problem, and nudges help them to overcome that problem; and (3) those in which choosers would be content with the outcomes produced by two or more nudges, or in which ex post preferences are endogenous to or constructed by nudges, so that the AJBT criterion leaves choice architects with several options, without specifying which one to choose. Cases that fall in category (1) plainly satisfy the AJBT criterion, and there are many such cases. From the standpoint of the AJBT criterion, cases that fall in category (2) are also unobjectionable, indeed they can be seen as a subset of category (1); and they are plentiful. Cases that fall in category (3) create special challenges, which may lead us to make direct inquiries into people's welfare or to explore what informed, active choosers typically select.

An important question, pressed by Robert Sugden (Sugden, 2016), is simple: "Do people really want to be nudged towards healthy lifestyles?" That is an empirical question, and we have a great deal of evidence about it. The answer is "yes," at least in the sense that in numerous nations – including the United States, the United Kingdom, Germany, France, Italy, and Australia – strong majorities endorse nudges toward healthy lifestyles (Jung and Mellers 2016; Sunstein 2016; Reisch and Sunstein 2016; Sunstein, Reisch, and Rauber 2017). One might object that general attitudes toward nudges do not specifically answer Sugden's question, and that it is best not to ask people generally (1) whether they approve of nudges, but instead to ask more specifically (2) whether they themselves would like to be nudged (on potential differences between answers to the two questions, see Cornwell and Krantz 2014). But general approval of health-related nudges strongly suggests that the answer to (2) is probably "yes" as well, and in any case, the existing evidence suggests that the answer to (2), asked specifically, is also "yes" (Jung and Mellers 2016). To be sure, more remains to be learned on these issues.

These findings cannot speak to conceptual and normative questions. What does it even mean to say that people want to make the choices toward which they are being nudged? What is the relationship between people's preferences and imaginable nudges? In light of behavioral findings, how confidently can we speak of people's "preferences"? If people want to be nudged, why are they not already doing what they would be nudged to do? Those are important questions (and they do have empirical features).

Thaler and I are interested in "libertarian paternalism," in the form of approaches that preserve freedom of choice, but that steer people in a direction that will promote their welfare. Coercive paternalism is very different (Conly 2014). Reminders, warnings, information disclosure, invocations of social norms, and default rules are examples of what we have in mind (Sunstein 2016). Contrary to a common misconception, the AJBT criterion is not at all designed

to counter the charge of paternalism. On the contrary, Thaler and I explicitly embrace (a mild form of) paternalism, and the purpose of the criterion is *to discipline the content of paternalistic interventions*. Consider, for example, a GPS device. It will identify the best route for you, given the direction that you identify; it makes you better off by your own lights. It increases *navigability*. At the same time, it is paternalistic in the sense that it purports to know, better than you do, how to get where you want to go. Its paternalism is one of means, not ends (Sunstein 2014). Means paternalism is unquestionably a form of paternalism, but it will typically satisfy the AJBT criterion, because it is respectful of people's ends. To be sure, we might find cases in which the distinction between means and ends is not straightforward (*ibid.*), but most cases are easy.

It is not possible to understand the operation of the AJBT criterion without reference to examples. In countless cases, we can fairly say that *given people's antecedent preferences*, a nudge will make choosers better off AJBT. For example:

1. *Luke has heart disease, and he needs to take various medications. He wants to do so, but he is sometimes forgetful. His doctor sends him periodic text messages. As a result, he takes the medications. He is very glad to receive those messages.*
2. *Meredith has a mild weight problem. She is aware of that fact, and while she does not suffer serious issues of self-control, and does not want to stop eating the foods that she enjoys, she does seek to lose weight. Because of a new law, many restaurants in her city have clear calorie labels, informing her of the caloric content of various options. As a result, she sometimes chooses low-calorie offerings – which she would not do if she were not informed. She is losing weight. She is very glad to see those calorie labels.*
3. *Edna is a professor at a large university, which has long offered its employees the option to sign up for a retirement plan. Edna believes that signing up would be a terrific idea, but she has not gotten around to it. She is somewhat embarrassed about that. Last year, the university switched to an automatic enrollment plan, by which employees are defaulted into the university's plan. They are allowed to opt out, but Edna does not. She is very glad that she has been automatically enrolled in the plan.*

There is nothing unfamiliar about these cases. On the contrary, they capture a great deal of the real-world terrain of nudging, both by governments and by the private sector (Sunstein 2016; Halpern 2015). Choosers have a goal, or an assortment of goals, and the relevant choice architecture can make it easier or harder for them to achieve it or them. Insofar as we understand the AJBT criterion by reference to people's antecedent preferences, that criterion is met. Note that it would be easy to design variations on these cases in which nudges *failed* that criterion, because they would make people worse off by their own lights.

We could complicate the cases of Luke, Meredith, and Edna by assuming that they have clear antecedent preferences, that the nudge is inconsistent with those preferences, but that as a result of the nudge, their preferences are changed. For example:

Jonathan likes talking on his cell phone while driving. He talks to friends on his commute to work, and he does business as well. As a result of a set of vivid warnings, he has stopped. He is glad. He cannot imagine why anyone would talk on a cell phone while driving. In his view, that is too dangerous.

After the nudge, Luke, Meredith, Edna, and Jonathan believe themselves to be better off. Cases of this kind raise the question whether the AJBT criterion requires reference to ex ante or ex post preferences. That is a good question, which might be answered by making direct inquiries into people's welfare; I will turn to that question below. My main point is that as originally given, the cases of Luke, Elizabeth, and Edna are straightforward. Such cases are common (Halpern 2015; Thaler 2015).

Some cases can be seen as different, because they raise questions about self-control (or akrasia):

- 1. Ted smokes cigarettes. He wishes that he had not started, but he has been unable to quit. His government has recently imposed a new requirement, which is that cigarette packages must be accompanied with graphic images, showing people with serious health problems, including lung cancer. Ted is deeply affected by those images; he cannot bear to see them. He quits, and he is glad.*
- 2. Joan is a student at a large university. She drinks a lot of alcohol. She enjoys it, but not that much, and she is worried that her drinking is impairing her performance and her health. She says that she would like to scale back, but for reasons that she does not entirely understand, she has found it difficult to do so. Her university recently embarked on an educational campaign to reduce drinking on campus, in which it (accurately) notes that four out of five students drink only twice a month or less. Informed of the social norm, Joan finally resolves to cut back her drinking. She does, and she is glad.*

In these cases, the chooser suffers from a self-control problem and is fully aware of that fact. Ted and Joan can be seen as both planners, with second-order preferences, and doers, with first-order preferences (Thaler and Sunstein 2008; Thaler 2015). A nudge helps to strengthen the hand of the planner. It is possible to raise interesting philosophical and economic questions about akrasia and planner-doer models (Stroud and Tappolet 2003), but insofar as Ted and Joan welcome the relevant nudges, and do so ex ante as well as ex post, the AJBT criterion is met. In a sense, self-control problems require their own GPS devices and so can be seen to involve navigability; but for choosers who face such problems, the underlying challenge is qualitatively distinctive, and they recognize that fact.

In such cases, the AJBT criterion is met. But do people will acknowledge that they face a self-control problem? That is an empirical question, of course, and my own preliminary research suggests that the answer is "yes." On Amazon's Mechanical Turk, I asked about 200 people this question:

Many people believe that they have an issue, whether large or small, of self-control. They may eat too much, they may smoke, they may drink too much, they may not save enough money. Do you believe that you have any issue of self-control?

A whopping 70 percent said that they did (55 percent said “somewhat agree,” while 15 percent said “strongly agree”). Only 22 percent disagreed. (Eight percent were neutral.)

This is a preliminary test, of course. Whatever majorities say, the cases of Ted and Mary capture a lot of the territory of human life, as reflected in the immense popularity of programs designed to help people to combat addiction to tobacco (Halpern et al. 2015) and alcohol. We should agree that nudges that do the work of such programs, or that are used in such programs (ibid.), are likely to satisfy the AJBT criterion.

There are harder cases. In some of them, it is not clear if people have antecedent preferences at all. In others – as in the case of Jonathan -- their ex post preferences are an artifact of, or constructed by, the nudge. Sometimes these two factors are combined (as marketers are well aware). As Amos Tversky and Richard Thaler put it long ago, “values or preferences are commonly constructed in the process of elicitation” (Tversky and Thaler 1990). If so, how ought the AJBT criterion to be understood and applied?

For example:

- 1. George cares about the environment, but he also cares about money. He currently receives his electricity from coal; he knows that coal is not exactly good for the environment, but it is cheap, and he does not bother to switch to wind, which would be slightly more expensive. He is quite content with the current situation. Last month, his government imposed an automatic enrollment rule on electricity providers: People will receive energy from wind, and pay a slight premium, unless they choose to switch. George does not bother to switch. He says that he likes the current situation of automatic enrollment. He approves of the policy and he approves of his own enrollment.*
- 2. Mary is automatically enrolled in a Bronze Health Care Plan – it is less expensive than Silver and Gold, but it is also less comprehensive in its coverage, and it has a higher deductible. Mary prefers Bronze and has no interest in switching. In a parallel world (a lot like ours, but not quite identical, Wolf 1990), Mary is automatically enrolled in a Gold Health Care Plan – it is more expensive than Silver and Bronze, but it is also more comprehensive in its coverage, and it has a lower deductible. Mary prefers Gold and has no interest in switching.*
- 3. Thomas has a serious illness. The question is whether he should have an operation, which is accompanied with potential benefits and potential risks. Reading about the operation online, Thomas is not sure whether he should go ahead with it. Thomas’ doctor advises him to have the operation, emphasizing how much he has to lose if he does not. He decides to follow the advice. In a parallel world (a lot like ours, but not*

quite identical), Thomas's doctor advises him not to have the operation, emphasizing how much he has to lose if he does. He decides to follow the advice.

In the latter two cases, Mary and Thomas appear to lack an antecedent preference; what they prefer is an artifact of the default rule (in the case of Mary) or the framing (in the case of Thomas). George's case is less clear, because he might be taken to have an antecedent preference in favor of green energy, but we could easily understand the narrative to mean that his preference, like that of Mary and Thomas, is endogenous to the default rule.

These are the situations on which I am now focusing: People lack an antecedent preference, and what they like is a product of the nudge. Their preference is constructed by it. After being nudged, they will be happy and possibly grateful. We have also seen that even if people have an antecedent preference, the nudge might change it, so that they will be happy and possibly grateful even if they did not want to be nudged in advance.

In all of these cases, application of the AJBT criterion is less simple. Choice architects cannot contend that they are merely vindicating choosers' ex ante preferences. If we look ex post, people do think that they are better off, and in that sense the criterion is met. For use of the AJBT criterion, the challenge is that *however Mary and Thomas are nudged, they will agree that they are better off*. In my view, there is no escaping at least some kind of welfarist analysis in choosing between the two worlds in the cases of Mary and Thomas. There is a large question about which nudge to choose in such cases (for relevant discussion, see Dolan 2014). Nonetheless, the AJBT criterion remains relevant in the sense that it constrains what choice architects can do, even if it does not specify a unique outcome (as it does in cases in which people have clear ex ante preferences and in which the nudge does not alter them).

Again: Thaler and I embrace paternalism, and so the AJBT criterion is emphatically not designed to defeat a charge of paternalism. It is psychologically fine (often) to think that choosers have antecedent preferences (whether or not "latent"), but that because of a lack of information or a behavioral bias, their choices will not satisfy them. (Recall the cases of Luke, Meredith, and Edna.) To be sure, it is imaginable that some forms of choice architecture will affect people who have information or lack such biases; an error-free cafeteria visitor might grab the first item she sees, because she is busy, and because it is not worth it to her to decide which item to choose (Ullmann-Margalit 2017). Consider this case:

Gretchen enjoys her employer's cafeteria. She tends to eat high-calorie meals, but she knows that, and she likes them a lot. Her employer recently redesigned the cafeteria so that salads and fruits are the most visible and accessible. She now chooses salad and fruit, and she likes them a lot.

By stipulation, Gretchen suffers from no behavioral bias, but she is affected by the nudge. But in many (standard) cases, behaviorally biased or uninformed choosers will be affected by a nudge, and less biased and informed choosers will not; a developing literature explores how to proceed in such cases, with careful reference to what seems to me a version of the AJBT criterion (Goldin 2015; Goldin and Lawson 2016).

In Gretchen's case, and all those like it, the criterion does not leave choice architects at sea: If she did not like the salad, the criterion would be violated. From the normative standpoint, it may not be entirely comforting to say that nudges satisfy the AJBT criterion if choice architects succeed in altering the preferences of those whom they are targeting. (Is that a road to serfdom? Recall the chilling last lines of Orwell's 1984: "He had won the victory over himself. He loved Big Brother" (cf. Elster 1983).) But insofar as we are concerned with subjective welfare, it is a highly relevant question whether choosers believe, *ex post*, that the nudge has produced an outcome of which they approve.

Countless nudges increase navigability, writ large, in the sense that they enable people to get where they want to go, and therefore enable them to satisfy their antecedent preferences. Many other nudges, helping to overcome self-control problems, are warmly welcomed by choosers, and so are consistent with the AJBT criterion. Numerous people acknowledge that they suffer from such problems. When people lack antecedent preferences or when those preferences are not firm, and when a nudge constructs or alters their preferences, the AJBT criterion is more difficult to operationalize, and it may not lead to a unique solution. But it restricts the universe of candidate solutions, and in that sense helps to orient choice architects.

References

- Conly, S. (2014). *Against Autonomy*. Oxford University Press.
- Cornwell, J. F., & Krantz, D. H. (2014). Public policy for thee, but not for me: Varying the grammatical person of public policy justifications influences their support. *Judgment and Decision Making*, 5, 433–444.
- Dolan, P. 2014. *Happiness By Design*. Penguin.
- Elster, J. 1983. *Sour Grapes*. Cambridge University Press.
- Goldin, J. (2015). Which Way to Nudge? Uncovering Preferences in the Behavioral Age. *Yale Law Journal*. 125, 226-271.
- Goldin, J. and Lawson, N. (2016). Defaults, Mandates, and Taxes: Policy Design with Active and Passive Decision-Makers. *American Journal of Law and Economics*. 18, 438-462.
- Halpern, S.D. et al. (2015). Randomized Trial of Four Financial-Incentive Programs for Smoking Cessation. *The New England Journal of Medicine*. 372:2108-211.
- Halpern, D. (2015) *Inside the Nudge Unit: How Small Changes Can Make a Big Difference*, W.H. Allen.
- Infante G, Lecouteux G, Sugden R (2016) Preference purification and the inner rational agent: a critique of the conventional wisdom of behavioural welfare economics. *J Econ Methodol* 23:1–25
- Jung JY, Mellers BA (2016) American Attitudes Toward Nudges. *Judgment and Decision Making* 11(1), 62–74.
- Lichtenstein, S. and Slovic, P. 2006. *The Construction of Preference*.
- Reisch, L. and Sunstein, C.R. Do Europeans Like Nudges? *Judgment and Decision Making*. 11: 310-25.
- Sugden, R. 2016. Do People Really Want to Be Nudged Towards Healthy Lifestyles? *International Review of Economics*.
- Stroud, S., and Tappolet, C. (eds.), 2003. *Weakness of Will and Practical Irrationality*, Oxford: Clarendon Press.
- Sunstein, C.R., Reich, L., and Rauber, J. (2017). Behavioral Insights All Over the World? Public Attitudes Toward Nudging in a Multi-Country Study, available at

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2921217

Sunstein, C.R. (2016) The Council of Psychological Advisers. *Annual Review of Psychol.* 67, 713–737.

Sunstein, C.R. (2016). *The Ethics of Influence*. Cambridge University Press.

Sunstein, C.R. 2014. *Why Nudge?* Yale University Press.

Sunstein, C.R., and Reisch, L.A. (2014). Automatically green: behavioral economics and environmental protection. *Harvard Environmental Law Review*. 38, 127-158.

Thaler, R.H. and Sunstein, C.R. (2008) *Nudge: Improving Decisions about Health, Wealth, and Happiness*, Yale University Press.

Thaler, R.H. (2015). *Misbehaving*. Norton.

Tversky, A. and Thaler, R.H. 1990. Anomalies: Preference Reversals. *Journal of Economic Perspectives*. 201-11.

Ullmann-Margalit, E. (2017). *Normal Rationality*. Oxford University Press.

Wolf, F. (1990). *Parallel Universes: The Search for Other Worlds*. Simon & Schuster.