

Decisional enhancement and autonomy: public attitudes towards overt and covert nudges

Gidon Felsen*

Noah Castelo[†]

Peter B. Reiner[‡]

Abstract

Ubiquitous cognitive biases hinder optimal decision making. Recent calls to assist decision makers in mitigating these biases—via interventions commonly called “nudges”—have been criticized as infringing upon individual autonomy. We tested the hypothesis that such “decisional enhancement” programs that target overt decision making—i.e., conscious, higher-order cognitive processes—would be more acceptable than similar programs that affect covert decision making—i.e., subconscious, lower-order processes. We presented respondents with vignettes in which they chose between an option that included a decisional enhancement program and a neutral option. In order to assess preferences for overt or covert decisional enhancement, we used the contrastive vignette technique in which different groups of respondents were presented with one of a pair of vignettes that targeted either conscious or subconscious processes. Other than the nature of the decisional enhancement, the vignettes were identical, allowing us to isolate the influence of the type of decisional enhancement on preferences. Overall, we found support for the hypothesis that people prefer conscious decisional enhancement. Further, respondents who perceived the influence of the program as more conscious than subconscious reported that their decisions under the program would be more “authentic”. However, this relative favorability was somewhat contingent upon context. We discuss our results with respect to the implementation and ethics of decisional enhancement.

Keywords: behavioral economics and policy, decision aids, contrastive vignettes.

1 Introduction

It has been well established that humans do not behave as fully rational actors but instead exhibit pervasive and predictable biases in decision-making (Ariely, 2008; Tversky & Kahneman, 1974). Efforts to counteract these biases in order to improve the outcome of decisions could be considered akin to, but distinct from, traditional cognitive enhancement (Farah et al., 2004) in that the goal is maximization of the value of choice outcomes, rather than improvement in particular domains of cognitive function that contribute to intelligence.¹ Recent proposals for such “decisional enhancement” have attracted interest from such diverse disciplines as public policy, psychology, economics, law, and ethics (Jolls et

al., 1998; Camerer et al., 2003; Rachlinski, 2002; Sunstein & Thaler, 2003; Trout, 2005; Blumenthal-Barby & Burroughs, 2012; Bovens, 2009). One well-known idea proposes modifying the environment in which choices are made such that individuals are “nudged” into making better decisions (Thaler & Sunstein, 2008). For example, requiring employees to opt out of, rather than opt in to, a retirement savings program takes advantage of the natural tendency to select the default option, thereby promoting the decision to save more for retirement. This approach has been dubbed “libertarian paternalism” in that a key feature of such decisional enhancement programs is that, while certain choices become more likely than others, the programs do not restrict the range of choices available to the individual. Although more employees will choose to save under the program described above, any individual employee remains free to choose to save any amount they want, including nothing.

Despite this attempt to preserve the range of available options, a key ethical issue engendered by the intentional attempt to influence individuals’ choices—even for their own benefit—is the potential for infringement on individual autonomy (Blumenthal-Barby & Burroughs, 2012; Dworkin, 1988; Hill, 2007). One line of criticism suggests that deigning to know which choice is best for the individual based on presumed objective measures of the choice outcome disregards other, unknowable, factors

This study was supported by grants from the Boettcher Foundation (G.F.) and the Canadian Institutes of Health Research (P.B.R.). The authors thank Adrian Carter for helpful comments on the manuscript.

Copyright: © 2013. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

*Department of Physiology & Biophysics, and Neuroscience Program, University of Colorado School of Medicine, 12800 E. 19th Ave., Mail Stop 8307, Aurora, CO 80045. Email: gidon.felsen@ucdenver.edu.

[†]Psychology Department, University of Toronto.

[‡]National Core for Neuroethics, University of British Columbia, Vancouver, BC, Canada

¹Note that increasing intelligence alone does not defend against maladaptive decision biases (Stanovich & West, 2008).

that contribute to the decision, such as the individual's value system (White, 2008). For example, even if participating in a retirement savings program would objectively increase an employee's future wealth she may oppose the program because it is inconsistent with her higher-order beliefs about investing, and therefore choosing to not participate in the program should not be considered a "poor" decision. According to this line of criticism, manipulating how the options are presented with the express purpose of making her more likely to join the program thus subverts her belief system, which disrespects her dignity as a person capable of autonomous decisions.

Thus, there is an inevitable tension between the respect for autonomy and the potentially beneficial outcomes—for both the individual and society—of improved decisions (Blumenthal-Barby & Burroughs, 2012). It is therefore of interest to determine how the public—key stakeholders in any social engineering program—deems the acceptability of decisional enhancement, and whether the degree of acceptability is dependent on particular features of the program (Castelo et al., 2012; House of Lords Science and Technology Select Committee, 2011). Although some studies have surveyed the levels of acceptance of generic decisional enhancement programs across different populations (Branson et al., 2011), few have examined the factors of specific programs that may affect acceptability (Marteau et al., 2011). Our study explores the question of whether the nature of the influence on decisions—specifically, whether it is overt or covert, a dimension particularly relevant for the consideration of autonomy—affects public acceptability of decisional enhancement programs.

In principle, decisional enhancement could affect one or more of the processes involved in making autonomous decisions. Three criteria are generally accepted as being required for a decision to be considered autonomous: it must be (1) consistent with an individual's conscious, higher-order desires; (2) rational, made with sufficient time and information to allow reflection; and (3) free from covert external influence (Christman, 1991; Dworkin, 1976; Dworkin, 1988; Felsen & Reiner, 2011; Frankfurt, 1971; Sugden, 1991; Taylor, 2005). Overt influences on decisions—i.e., those of which the decision maker is aware and can consciously process—do not appear to violate any of these conditions. Covert influences, however, present a challenge. While they plainly violate the third condition and therefore infringe upon autonomy, they may also be seen to *enhance* autonomy by increasing the likelihood of a decision aligned with the first condition, their higher-order desires (Trout, 2005). For example, subconsciously decreasing hunger (a lower-order desire) would make refraining from overeating (a higher-order desire) more likely.

While some studies suggest that covert influences may

have a greater effect than overt influences on choice outcome (e.g., Duffy & Verges, 2008; Goldstein et al., 2008; Nolan et al., 2008; Wisdom et al., 2010), to our knowledge the question of which is more acceptable has not been addressed empirically. If covert influences are found to be more acceptable, then the most effective decisional enhancement programs would also be the most acceptable, providing a clear path for implementation of such programs. If, however, covert influences are found to be less acceptable, then the most effective and most acceptable strategies would be in conflict and policy makers would need to consider not only the value of improved decision making but also the tradeoff in terms of individual autonomy.

We hypothesized that overt influences would be more acceptable than covert influences. Consistent with the multiple necessary conditions for autonomy and the complexity of human decision processes, we found that our hypothesis was supported within most contexts (eating, purchasing, exercising, and investing decisions), but not in another (workplace productivity decisions). We discuss our findings within the framework of autonomous decision making, with an eye towards their application to the policy debate on, and ethics of, decisional enhancement.

2 Method

We used a between-subjects design—the contrastive vignette technique (Burstin et al., 1980)—to probe public attitudes towards manipulations that employed either conscious (overt) or subconscious (covert) decisional enhancement. Different versions of a single vignette with minimal variations were presented (Table 1), and respondents across conditions (subconscious and conscious) answered an identical set of questions (Table 2). Following acceptance of informed consent and completion of a brief set of demographic questions, each respondent was randomly assigned to see one and only one version of the vignette, within only one context (e.g., healthy eating), and was blind to the contrastive condition. The effects of the type of influence on decisions (e.g., subconscious), can be measured by comparing responses across conditions. Carefully constructed contrastive vignettes therefore allow us to control for demand characteristics, which are features of the experiment that could potentially alert participants to the study's hypothesis. Such demand characteristics could in turn cause participants to alter their responses, consciously or not, in order to support or undermine the hypothesis or according to social desirability concerns (Orne, 1962; Nichols & Maner, 2008).

Respondents from the United States and Canada were recruited via Amazon's Mechanical Turk (Ipeirotis, 2010);

Table 1: Questions in the Eating scenario. All respondents within each scenario were asked the same 5 questions (in addition to the comprehension check). The wording of some questions varied slightly across scenarios as appropriate; shown here are the questions asked in the Eating scenario. The vignette was presented following Question 1. Anchors for each 9-point Likert scale, and their corresponding numerical values (as superscript), are shown following each question.

- Q1:** To what extent do you feel like you could use help making healthier eating choices in the face of the availability of unhealthy but tasty foods?
[Not at all¹ / Very much⁹]
[Presentation of vignette]
- Q2:** How would the existence of this program affect the likelihood that you would accept the job offer from Company B, the company that offers the program that encourages healthy eating?
[Much less likely¹ / Much more likely⁹]
- Q3:** If you did accept the job offer at Company B, to what extent do you think your decisions about food choices in the Company cafeteria would reflect your authentic preferences?
[Not at all reflect my authentic preferences¹ / Completely reflect my authentic preferences⁹]
- Q4:** How do you think that the program described on the previous page affects decision making? Please answer using the scale below, where:
“1” represents “Entirely subconsciously”, and
“9” represents “Entirely consciously”.
[Entirely subconsciously¹ / Entirely consciously⁹]
- Q5:** (Optional) If you have time, please tell us why you answered as you did.

Paolacci et al., 2010). Respondents were compensated \$0.25 for completion of the survey. Once they accepted the assignment, they were directed to an external website, where they were randomly assigned to one of ten vignettes, composed of five pairs of contrastive vignettes.

The two versions of each pair of contrastive vignettes were designed to be as similar as possible in every respect except for the *type* of decisional influence: one purported to influence conscious processing and the other purported to influence subconscious processing. Differences in responses to the paired contrastive vignettes can therefore be largely attributed to whether the manipulation was conscious or subconscious. Five scenarios were explored: healthy eating, prudent purchasing, increased exercise, prudent investing, and productivity at work (referred to below as the Eating, Purchasing, Exercising, Investing,

and Productivity scenarios, respectively). For example, in the Eating scenario, respondents in the SUBCONSCIOUS group read that a “cafeteria has been revamped so that unhealthy foods, such as candy bars, potato chips, and the like are not as conveniently located”, while respondents in the CONSCIOUS group read that a “cafeteria has been revamped so that all foods have their nutritional content clearly displayed” (Table 1). The scenarios are in the Appendix. The effectiveness at reducing caloric intake of manipulations along each of these lines has been studied (Chapman & Ogden, 2012; Downs et al., 2009; Harnack & French, 2008; Pulos & Leng, 2010; Roberto et al., 2010; Tandon et al., 2011; Wisdom et al., 2010), although to our knowledge their relative acceptability has not been examined.

Prior to the presentation of the vignettes, respondents were asked how much they felt they could use help making decisions within the context of the assigned scenario (Table 1, Question 1). This introductory query was used to determine whether attitudes towards subconscious and conscious decisional enhancement depended upon a desire for assistance.

Following each vignette, respondents were asked to consider a choice between two potential employers (Eating, Exercise, Investing, and Productivity) or credit card companies (Purchasing), where one employer/company offered the relevant decisional enhancement program while the other did not (the “neutral option”) (Table 1, Question 2). The primary outcome measure was the rating on a 9-point scale of the degree to which the program would affect the likelihood of favoring the option with the decisional enhancement program over the neutral option.

Three follow-up questions explored various aspects of respondents’ perceptions of the vignettes. First, respondents were asked to rate how authentic their decisions would be if they decided to select the option with the decisional enhancement program (Table 1, Question 3). Next, respondents were asked how the previously described program affected decision making (Table 1, Question 4). We used these responses as a measure of how the respondents *perceived* the effect of the program on decision making, which for some analyses is more informative than the stated condition (subconscious or conscious). Respondents were then offered the option of completing a free-form text box which asked them to tell us why they answered as they did (Table 1, Question 5). Finally, comprehension of the vignettes was verified; for example, respondents in the Eating scenario were asked which of the following “The previous question discussed”: “Strength training”, “Healthy food choices” (the correct answer), “Prudent online purchasing”, and “Improving productivity”. Correctly answering the comprehension check was required for responses to be included in the data set. The full set of vignettes, and questions

presented for one representative scenario, can be found in the Appendix and Table 1, respectively.

3 Results

We collected data from 2,775 respondents from the United States and Canada who correctly answered the comprehension question at the end of the survey (mean age: 29.8 years; 40% female). Since our primary objective was to explore whether respondents felt that decisional enhancement programs that overtly affect conscious reasoning were more acceptable than those that covertly affect subconscious thought processes, we examined respondents' stated perceptions of the manipulation (as affecting either conscious or subconscious decision making) to ensure the validity of our analyses. We found that, despite our explicit description of the manipulations as affecting either conscious or subconscious processing (Table 1), a sizeable fraction of each group perceived the intervention as incongruent with our description: 50% of the SUBCONSCIOUS group rated the manipulation as more conscious than subconscious, and 17% of the CONSCIOUS group rated the manipulation as more subconscious than conscious. In order to first analyze only the data from respondents whose perception matched our intended manipulation, we excluded respondents whose perception of the manipulation was incongruent with the assigned condition. We then tested the hypothesis that congruent respondents in the CONSCIOUS group would be more likely than congruent respondents in the SUBCONSCIOUS group to favor the option with the decisional enhancement program over the neutral option. The results support this hypothesis in the Eating, Purchasing, Exercising, and Investing scenarios, but not in the Productivity scenario (Figure 1; Eating, $p = 0.0064$; Purchasing, $p \sim .000$; Exercising, $p = 0.004$; Investing, $p \sim .000$; Productivity, $p = 0.43$). Results were similar when all respondents were included: Eating, $p = 0.38$; Purchasing, $p = 0.004$; Exercising, $p = 0.025$; Investing, $p \sim .000$; Productivity, $p = 0.83$; 1-tailed t-tests).


In order to corroborate these results, we examined whether the likelihood of favoring the option with the decisional enhancement program (Table 1, Question 2) correlated with the respondents' stated perception of how the program affected decision making (Table 1, Question 4). We included all respondents in this analysis (and in all subsequent analyses below), because we were specifically interested in how their stated perception of the manipulation as conscious or subconscious correlated with their stated preference for the option with the decisional enhancement program. In agreement with the results described above (Figure 1), there was a significant positive correlation between these variables in all scenarios except

for the Productivity scenario (Table 2, row 1). In further support of our hypothesis, we found that respondents' stated perception of the manipulation correlated with the extent to which they favored the option with the decisional enhancement program even within their assigned group (Table 2, rows 2-3).

Any universally applied program—be it enacted by government, corporations, or other large organizations—will necessarily affect individuals who do not want help as well those who do. We therefore examined how the sentiments of respondents who reported that they could use help making decisions compared to those of respondents who reported that they could not use help. As expected, across all scenarios respondents in both the SUBCONSCIOUS and CONSCIOUS groups who wanted help were more likely to favor the option with the decisional enhancement program over the neutral option (Table 2, rows 4-5). However, the degree to which respondents wanted help had little effect on the relative favorability for the option with the conscious rather than subconscious decisional enhancement program: Only in the Eating scenario did this variable moderate the size of the effect (Table 2, row 6). Thus, we found support for our hypothesis that decisional enhancement that affects conscious deliberation is more acceptable than decisional enhancement that affects subconscious processes, but our results did not generalize across all five scenarios. We discuss possible interpretations of these results below.

Finally, we examined whether the type of manipulation (conscious or subconscious) affected the perceived authenticity of respondents' decisions within the context of the decisional enhancement program. We first examined whether stated perceptions of how the program affected decision making (Table 2, Question 4) correlated with the extent to which decisions in the context of the program would be perceived as reflecting authentic preferences (Table 2, Question 3). We found a significant positive correlation in all scenarios (Table 3, row 7). In addition, we found that respondents' likelihood of opting for the decisional enhancement program was related to whether the decisions they made were perceived as being authentic: In all scenarios, there was a significant positive correlation between the degree to which respondents felt that decisions made within the context of the decisional enhancement program would be authentic and how likely they were to opt for the program (Table 3, row 8; Figure 2). These data support the idea that preserving the individual's capacity for making authentic decisions is an important condition for the acceptability of decisional enhancement programs.

Figure 1: Effect of condition (subconscious or conscious influence) on the relative favorability of the option with the decisional enhancement program over the neutral option. 1-tailed t-tests examined whether the CONSCIOUS group was more likely than the SUBCONSCIOUS group to favor the option with the decisional enhancement program, in each scenario. Respondents who perceived the effect of the program on decisions incongruently with their assigned condition were excluded. *, $p < 0.01$. Error bars, \pm SEM.



4 Discussion

Progress in the behavioral sciences has revealed the many ways in which human decision-making predictably departs from the rational actor model (Ariely, 2008; Kahneman, 2011; Stanovich & West, 2000; Strack & Deutsch, 2004; Tversky & Kahneman, 1974). This insight has inspired “choice architects”, i.e., engineers of the environments in which decisions are made, to influence decisions according to their desired outcomes (which may or may not align with the desired outcomes of the decision makers). As more organizations consider implementing such programs, debate has arisen about the acceptability of their use given the frequent lack of objective criteria about what constitutes a better decision outcome, as well as the desire to protect autonomous decision making (Blumenthal-Barby & Burroughs, 2012; Bovens, 2009; Hausman & Welch, 2010; Hill, 2007; Mitchell, 2004; House of Lords Science and Technology Select Committee, 2011; Sugden, 2008; Sunstein & Thaler, 2003; White, 2008). The dominant narrative in this debate assumes that decisional enhancement necessarily infringes upon individual autonomy, and proceeds to weigh this cost against the benefits of the manipulation. There are, however, theoretical reasons to question the generality of this tradeoff. As described in the introduction, covert influences on decisions may subvert the autonomy of the decision maker. However, covertly influencing decision processes such that the resulting decision is *aligned with higher-order* desires may actually enhance autonomy (Trout, 2005), especially in situations in which the target population is known to want help with a given behavior. How these considerations interact in order to determine whether autonomy is reduced or en-

hanced by particular manipulations, and how this interaction depends on the context and the goals of the decision maker herself, are open questions that merit empirical study.

As an initial foray into empirically addressing the question of how to improve decision making while infringing minimally on autonomy, the present study examined public attitudes towards decisional enhancement programs intended to influence decision making either covertly or overtly. We probed how likely respondents would be to participate in such a program, as well as their perceptions of how authentic their resulting decisions would be. While overall we found support for our hypothesis that overt, rather than covert, influences would be more acceptable, our results depended to some extent on the specific context. In the Eating, Purchasing, Exercising, and Investing scenarios, but not the Productivity scenario, the relative favorability of the decisional enhancement program over the neutral option was higher when the influence on decision processes was conscious than when it was subconscious (Figure 1). In the Eating scenario only, the relative favorability of covert influences was moderated by the degree to which respondents wanted help such that, the less respondents wanted help, the more favorable they were to the conscious than the subconscious influence (Table 2, row 6). Under no conditions were the respondents more favorable to the subconscious influence than to the conscious influence. Further, we found that the degree to which respondents believed that their decisions within the context of the program would be authentic correlated with 1) the degree to which the influence affected conscious decision processes, and 2) how likely they were to favor the decisional enhancement program over the neutral option (Figure 2).

Table 2: Correlation and regression results. Data from all respondents are included. Q1-4 refer to Questions 1-4 (see Table 2).

Test	Eating		Purchasing		Exercise		Investing		Productivity	
	r	p	r	p	r	p	r	p	r	p
Correlation between responses to Q2 and Q4, all respondents	.23	.000	.18	.000	.12	.009	.28	.000	.08	.08
Correlation between responses to Q2 and Q4, subconscious	.21	.000	.13	.017	.06	.29	.19	.000	.10	.047
Correlation between responses to Q2 and Q4, conscious	.31	.000	.12	.001	.24	.004	.30	.000	.04	.55
Correlation between responses to Q1 and Q2, subconscious	.36	.000	.29	.000	.35	.000	.30	.000	.23	.000
Correlation between responses to Q1 and Q2, conscious	.16	.024	.31	.000	.32	.000	.23	.000	.25	.000
Regression of interaction between responses to Q1 and Q4 on Q2, all respondents (β)	-.33	.039	.02	.32	-.00	.93	.00	.93	.00	.84
Correlation between responses to Q3 and Q4, all respondents	.23	.000	.20	.000	.15	.000	.25	.000	.11	.011
Correlation between responses to Q2 and Q3, all respondents	.39	.000	.45	.000	.57	.000	.49	.000	.32	.000

Together, these results suggest that public acceptance of a given intervention may depend on the degree to which it infringes upon autonomy, but is also affected by other context-specific factors.

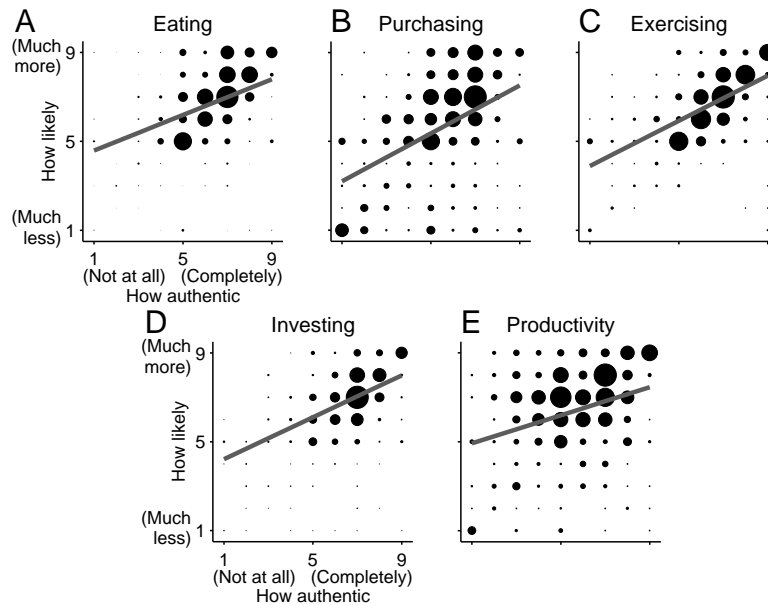
It is notable that we did not observe a preference for overt over covert influences in the Productivity scenario (Figure 1; Table 2, row 1). One possibility that we considered was that respondents perceived the outcome of “enhanced” decisions as not benefiting them personally, i.e., not being in their own best interest. In an initial version of this vignette, the benefit to the “consultant” (i.e., the respondent) was ambiguous and respondents were no more favorable to conscious than subconscious decisional enhancement (data not shown). For this reason, we modified our original Productivity vignettes to make it clear that the benefits of the program accrued to the “consultant” and not to the company (Appendix), but we still found that the CONSCIOUS group was no more likely than the SUBCONSCIOUS group to favor the option with the decisional enhancement program over the neutral option. Thus, the lack of an effect in the Productivity scenario is not due to a perceived absence of personal benefit from the decisional enhancement program.

The general observation that context has a modest but meaningful impact upon preference for conscious versus subconscious decisional enhancement should give choice architects pause. Several possible explanations are wor-

thy of further study. For example, it is possible that decisions in some contexts are seen as less “one’s own” in the first place, perhaps due to the personal experience of being unable to control their decisions, and therefore individuals may be more accepting of decisional enhancement programs in those contexts. Another possibility is that people may feel that decisions made in some contexts are more consciously driven than decisions made in other contexts. For example, respondents who wanted help with eating decisions may have been more likely to recognize that food choices are often subconsciously driven, and were therefore just as likely to favor the decisional enhancement program with covert influences as the program with overt influences, whereas respondents who did not want help with food choices reverted to the expected preference for overt influences (Table 3, row 6).

It is worth noting some limitations of this study. Most obviously, the vignettes are hypothetical: our respondents reported how they think they would act if placed in a particular situation, which may differ from how they would actually act in that real situation (Chang et al., 2009). A second limitation is the categorical distinction between conscious and subconscious manipulations. Each of the scenarios clearly described the manipulation as either explicitly “conscious” or “subconscious” (Appendix). Nevertheless, a substantial fraction of each group did not perceive the influence as explicitly stated, rating subcon-

Figure 2: Correlation between the authenticity of decisions made within the context of the program and the relative favorability of the option with the decisional enhancement program over the neutral option. (A) Paired responses to Questions 2 and 3 (see Table 2) are shown for all respondents in in the Eating scenario. Best-fit line shown in gray. (B-E) As in (A), in the Purchasing, Exercising, Investing, and Productivity scenarios, respectively. Circle size corresponds to number of respondents for each pair of responses, normalized within each scenario.



conscious influences as more conscious than subconscious, and vice versa. This incongruence may have been due to inattentive respondents, who failed to understand our descriptions, but these same respondents passed our comprehension check, suggesting that it was not due to lack of attention. Instead, their stated opinions about the influence may reflect the fact that no influence affects *only* conscious or subconscious processes. It would therefore be more accurate to consider the influence of a decisional enhancement program to lie somewhere along the continuum from covert to overt. Indeed, this is why we examined how stated perceptions of the degree to which the program affected conscious processing correlated with several variables (Table 2). Despite these limitations, the present study is a necessary first step towards grounding the debate surrounding autonomy and the use of decisional enhancement in empirical data. It would be useful for future studies to examine these issues in “real life”, as opposed to survey-based, situations.

Another possible limitation is that the use of contrastive vignettes does not inform us about attitudes toward public policies, where the policy question is explicitly about whether conscious or subconscious manipulation should be preferred. To examine this possibility, we collected data from a new set of respondents using a direct comparison test that was identical to our Eating scenario except that respondents chose between Company

A, which employed a decisional enhancement program in its cafeteria that targeted subconscious processes, and Company B, which employed a decisional enhancement program in its cafeteria that targeted subconscious processes. We found that respondents significantly favored Company B (with the program targeting conscious processing) ($p < 0.01$, one-tailed t-test, $n = 155$). The magnitude of this effect ($5.52 - 5.00 = 0.52$) was similar to that observed in the contrastive vignette version of the Eating scenario (Conscious condition: 6.81; Subconscious condition: 6.33; difference: $6.81 - 6.33 = 0.48$).

These data have clear implications for public policy. Proponents of decisional enhancement may hesitate to enact programs that enhance decision making if they see the influence as restricting the individual’s ability to choose freely—an effect that may be seen as particularly pernicious with covert influences (White, 2008). However, where public attitudes are indifferent between conscious and subconscious influences, we suggest that policy makers gain license to use the most effective tools at their disposal, even those that are covert. This recommendation is consistent with the possibility that the subconscious processing of covert external influences is sufficiently pervasive to call into question the degree to which many decisions can be autonomous at all (Felsen & Reiner, 2011). Hence, well-meaning attempts to “preserve” autonomy at the expense of improved decision

outcomes may be misguided. On the other hand, in situations where public attitudes do indicate a preference for overt influences, or a distaste for covert influences, due consideration to the balance between outcomes and the preservation of perceived autonomy should be informed by empirical data.

References

- Ariely, D. (2008). *Predictably irrational: The hidden forces that shape our decisions*. New York: Harper.
- Blumenthal-Barby, J. S., & Burroughs, H. (2012). Seeking better health care outcomes: the ethics of using the “nudge.” *The American Journal of Bioethics*, *12*, 1–10.
- Bovens, L. (2009). The ethics of nudge. In T. Grüne-Yanoff & S. O. Hansson (Eds.), *Preference Change: Approaches from Philosophy, Economics and Psychology* (pp. 207–220). Berlin and New York: Springer.
- Branson, C., Duffy, B., Perry, C., & Wellings, D. (2011). *Acceptable behaviour? Public opinion on behaviour change policy*. Ipsos Social Research Institute.
- Burstin, K., Doughtie, E. B., & Raphaeli, A. (1980). Contrastive vignette technique: An indirect methodology designed to address reactive social attitude measurement. *Journal of Applied Social Psychology*, *10*, 147–165.
- Camerer, C., Issacharoff, S., Loewenstein, G., O’Donoghue, T., & Rabin, M. (2003). Regulation for conservatives: Behavioral economics and the case for “asymmetric paternalism.” *University of Pennsylvania Law Review*, *151*, 1211–1254.
- Castelo, N., Reiner, P. B., & Felsen, G. (2012). Balancing autonomy and decisional enhancement: An evidence-based approach. *The American Journal of Bioethics*, *12*, 30–31.
- Chang, J. B., Lusk, J. L., & Norwood, F. B. (2009). How closely do hypothetical surveys and laboratory experiments predict field behavior? *American Journal of Agricultural Economics*, *91*, 518–534.
- Chapman, K., & Ogden, J. (2012). Nudging customers towards healthier choices: an intervention in the university canteen. *Journal of Food Research*, *1*, 13–21.
- Christman, J. (1991). Autonomy and personal history. *Canadian Journal of Philosophy*, *21*, 1–24.
- Downs, J. S., Loewenstein, G., & Wisdom, J. (2009). Strategies for promoting healthier food choices. *American Economic Review*, *99*, 159–164.
- Duffy, S., & Verges, M. (2008). It matters a hole lot: Perceptual affordances of waste containers influence recycling compliance. *Environment and Behavior*, *41*, 741–749.
- Dworkin, G. (1976). Autonomy and behavior control. *Hastings Center Report*, *6*, 23–28.
- Dworkin, G. (1988). *The theory and practice of autonomy*. Cambridge, UK: Cambridge University Press.
- Farah, M. J., Illes, J., Cook-Deegan, R., Gardner, H., Kandel, E., King, P., Parens, E., Sahakian, B., & Wolpe, P. R. (2004). Neurocognitive enhancement: What can we do and what should we do? *Nature Reviews Neuroscience*, *5*, 421–425.
- Felsen, G., & Reiner, P. B. (2011). How the neuroscience of decision making informs our conception of autonomy. *AJOB Neuroscience*, *2*, 3–14.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy*, *68*, 5–20.
- Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, *35*, 472–482.
- Harnack, L. J., & French, S. A. (2008). Effect of point-of-purchase calorie labeling on restaurant and cafeteria food choices: A review of the literature. *International Journal of Behavioral Nutrition and Physical Activity*, *5*, 51.
- Hausman, D. M., & Welch, B. (2010). Debate: To nudge or not to nudge. *Journal of Political Philosophy*, *18*, 123–136.
- Hill, C. (2007). Anti-Anti-Anti-Paternalism. *New York University Journal of Law and Liberty*, *2*, 444–454.
- Ipeirotis, P. G. (2010). Demographics of mechanical turk. Retrieved from <http://www.ipeirotis.com/wp-content/uploads/2012/02/CeDER-10-01.pdf>. Accessed 8/13/2012
- Jolls, C., Sunstein, C. R., & Thaler, R. (1998). A behavioral approach to law and economics. *Stanford Law Review*, *50*, 1471–1550.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus, and Giroux.
- Marteau, T. M., Ogilvie, D., Roland, M., Suhrcke, M., & Kelly, M. P. (2011). Judging nudging: Can nudging improve population health? *BMJ*, *342*, 263–265.
- Mitchell, G. (2004). Libertarian paternalism is an oxymoron. *Northwestern University Law Review*, *99*, 1245–1277.
- Nichols, A., & Maner, J. (2008). The Good-Subject Effect: Investigating Participant Demand Characteristics. *The Journal of General Psychology*, *135*, 151–165.
- Nolan, J. M., Schultz, P. W., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2008). Normative social influence is underdetected. *Personality and Social Psychology Bulletin*, *34*, 913–923.
- Orne, M. T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, *17*, 776–783.
- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judg-*

- ment and Decision Making*, 5, 411–419.
- Pulos, E., & Leng, K. (2010). Evaluation of a voluntary menu-labeling program in full-service restaurants. *American Journal of Public Health*, 100, 1035–1039.
- Rachlinski, J. J. (2002). The uncertain psychological case for paternalism. *Northwestern University Law Review*, 97, 1165–1225.
- Roberto, C. A., Larsen, P. D., Agnew, H., Baik, J., & Brownell, K. D. (2010). Evaluating the impact of menu labeling on food choices and intake. *American Journal of Public Health*, 100, 312–318.
- House of Lords Science and Technology Select Committee. (2011). *Behaviour Change - Second Report*.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: implications for the rationality debate? *The Behavioral and Brain Sciences*, 23, 645–665.
- Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology*, 94, 672–695.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8, 220–247.
- Sugden, R. (1991). Rational choice: A survey of contributions from economics and philosophy. *The Economic Journal*, 101, 751–785.
- Sugden, R. (2008). Why incoherent preferences do not justify paternalism. *Constitutional Political Economy*, 19, 226–248.
- Sunstein, C. R., & Thaler, R. H. (2003). Libertarian paternalism is not an oxymoron. *University of Chicago Law Review*, 70, 1159–1202.
- Tandon, P. S., Zhou, C., Chan, N. L., Lozano, P., Couch, S. C., Glanz, K., Krieger, J., & Saelens, B. E. (2011). The impact of menu labeling on fast-food purchases for children and parents. *American Journal of Preventive Medicine*, 41, 434–438.
- Taylor, J. (2005). *Personal autonomy: New essays on personal autonomy and its role in contemporary moral philosophy*. Cambridge, UK: Cambridge University Press.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, CT: Yale University Press.
- Trout, J. D. (2005). Paternalism and cognitive bias. *Law and Philosophy*, 24, 393–434.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science*, 185, 1124–1131.
- White, M. D. (2008). Behavioral law and economics: the assault on consent, will, and dignity. In G. Gaus, C. Favor, & J. Lamont (Eds.), *Essays on philosophy, politics and economics: Integration and common research projects* (pp. 201–224). Palo Alto, CA: Stanford University Press.
- Wisdom, J., Downs, J. S., & Loewenstein, G. (2010). Promoting healthy choices: information versus convenience. *American Economic Journal: Applied Economics*, 2, 164–178.

Appendix: Contrastive vignettes. Respondents were randomly assigned to 1 of the 5 scenarios (Eating, Purchasing, Exercising, Investing, or Productivity) and to 1 of the 2 conditions (subconscious or conscious). Respondents within each scenario were presented with identical text with the exception of the [Subconscious] and [Conscious] columns. Note that within these columns, only the gray text differed across conditions.

[Eating scenario:]

You have been offered a job by two different companies (let's call them Company A and Company B). Everything about the two offers is identical, with one exception: Company B has teamed up with a nutrition behavior planning organization that has helped redesign the company cafeteria in a manner that is designed to encourage healthy eating.

[Subconscious:]

The new design works like this—the cafeteria has been revamped so that unhealthy foods, such as candy bars, potato chips, and the like are not as conveniently located. You can still choose whichever foods you would like, but moving the location of the unhealthy food in the cafeteria results in a subconsciously-driven bias towards healthy eating choices; in other words, the decision to eat healthy foods is made more likely without the need for conscious deliberation. Studies have shown that implementing this policy leads to healthier eating habits.

When you inquire about similar programs at Company A, they tell you that they have no program like that. Given your financial situation, you are going to accept a job offer with one of the companies.

[Conscious:]

The new design works like this - the cafeteria has been revamped so that all foods have their nutritional content clearly displayed. You can still choose whichever foods you would like, but the nutritional information results in a consciously-driven bias towards healthy eating choices; in other words, the decision to eat healthy foods is made more likely as a result of conscious deliberation. Studies have shown that implementing this policy leads to healthier eating habits.

[Purchasing scenario:]

You have been offered an opportunity to open a credit card account by two different companies (let's call them Company A and Company B). Everything about the two offers is identical, with one exception: Company B has teamed up with a consumer behavior planning organization that offers you a free program designed to encourage prudent online purchasing.

[Subconscious:]

The program works like this—whenever you browse for items to purchase online, the program uses your purchasing history to create a unified “marketplace” webpage, listing all items available for purchase on a single page. Items that are less likely to be necessary are listed towards the bottom of the page, and you must scroll down to see them; items that are more likely to be necessary are listed near the top. You can still choose to buy whatever you would like, but the requirement to scroll down to the unnecessary items results in a subconsciously-driven bias towards necessary purchases; in other words, the decision to make prudent online purchases is made more likely without the need for conscious deliberation. Studies have shown that using this program leads to more prudent online purchases.

[Conscious:]

The program works like this—whenever you browse for items to purchase online, the program uses your purchasing history to create a unified “marketplace” webpage, listing all items available for purchase on a single page. Your current credit card balance and monthly budget are displayed along a sidebar of the marketplace page. You can still choose to buy whatever you would like, but the financial information results in a consciously-driven bias towards necessary purchases; in other words, the decision to make prudent online purchases is made more likely as a result of conscious deliberation. Studies have shown that using this program leads to more prudent online purchases.

When you inquire about similar programs at Company A, they tell you that they have no program like that. Given your financial situation, you are going to open a credit card account with one of the companies.

[Exercising scenario:]

You have been offered a job by two different companies (let's call them Company A and Company B). Everything about the two offers is identical, with one exception: Company B has teamed up with a behavior planning organization that has helped redesign the two-story building in a manner that is designed to encourage exercise.

[Subconscious:]

The new design works like this—the building has been revamped so that getting to the second floor is faster via the stairs than the elevator because the elevator has been slowed down. You can still choose to take the stairs or the elevator, but the time saved via the stairs results in a subconsciously-driven bias towards increased exercise; in other words, the decision to exercise is made more likely without the need for conscious deliberation. Studies have shown that implementing this policy leads to an increase in exercise.

[Conscious:]

The new design works like this—the building has been revamped so that the relative health benefits of taking the stairs instead of the elevator to the second floor are clearly displayed. You can still choose to take the stairs or the elevator, but the exercise information results in a consciously-driven bias towards increased exercise; in other words, the decision to exercise is made more likely as a result of conscious deliberation. Studies have shown that implementing this policy leads to an increase in exercise.

When you inquire about similar programs at Company A, they tell you that they have no program like that. Given your financial situation, you are going to accept a job offer with one of the companies.

[Investing scenario:]

You have been offered a job by two different companies (let's call them Company A and Company B). Everything about the two offers is identical, with one exception: Company B has teamed up with an investment planning organization that has helped redesign the retirement savings plan in a manner that is designed to encourage investing for retirement.

[Subconscious:]

The new design works like this—with every annual salary increase you are provided information in the form of a series of icons representing tropical beaches that shows how much extra leisure you are likely to be able to afford during your retirement by investing different percentages of your increased salary; larger investments now translate into more retirement savings later. You can still choose to keep the entire salary increase instead of investing it, but the information provided results in a subconsciously-driven bias towards investment; in other words, the decision to invest is made more likely as a result of subconscious deliberation. Studies have shown that implementing this policy leads to an increase in retirement savings.

[Conscious:]

The new design works like this—with every annual salary increase you are provided information in the form of a detailed table of your earnings that shows how much extra money you are likely to have during your retirement by investing different percentages of your increased salary; larger investments now translate into more retirement savings later. You can still choose to keep the entire salary increase instead of investing it, but the information provided results in a consciously-driven bias towards long-term investment; in other words, the decision to invest is made more likely as a result of conscious deliberation. Studies have shown that implementing this policy leads to an increase in retirement savings.

When you inquire about similar programs at Company A, they tell you that they have a standard retirement plan, but no program to encourage investment. Given your financial situation, you are going to accept a job offer with one of the companies.

[Productivity scenario:]

You have been offered a contract as an independent consultant, paid based on your rate of productivity, by two different companies (let's call them Company A and Company B). Everything about the two offers is identical, with one exception: Company B has teamed up with a productivity behavior planning organization that offers you a free software program designed to encourage improved productivity.

[Subconscious:]

The software works like this—whenever you are working on the project, the program slows down your Internet browsing speed for non-work related websites. You can still choose to view whatever content you would like, but the inconvenience of slower browsing results in a subconsciously-driven bias towards work-related content; in other words, the decision to spend time productively is made more likely without the need for conscious deliberation. Studies have shown that using this software leads to higher rates of productivity, and therefore higher pay as per the contract.

[Conscious:]

The software works like this—whenever you are working on the project, the program displays the time you spend on the Internet browsing non-work related websites, and alerts you when you exceed a certain time limit that you set for yourself. You can still choose to view whatever content you would like, but the time information results in a consciously-driven bias towards work-related content; in other words, the decision to spend time productively is made more likely as a result of conscious deliberation. Studies have shown that using this software leads to higher rates of productivity, and therefore higher pay as per the contract.

When you inquire about similar programs at Company A, they tell you that they have no program like that. Given your financial situation, you are going to accept a contract with one of the companies.
