NBER WORKING PAPER SERIES

NUDGES IN EXERCISE COMMITMENT CONTRACTS: A RANDOMIZED TRIAL

Jay Bhattacharya
Alan M. Garber
Jeremy D. Goldhaber-Fiebert

Working Paper 21406
http://www.nber.org/papers/w21406

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2015

The authors wish to thank students who aided in early parts of this study, Erik Blumenkrantz and Roxana Azimi. The authors are grateful to seminar participants at Stanford Health Policy's Research in Progress seminar as well as seminar participant the Universities of Oslo and of Bergen in Norway. We also thank Kevin Volpp and conference participants at University of Pennsylvania's annual meeting at the Center for Health Incentives and Behavioral Economics. We presented preliminary findings at the Society for Medical Decision Making's Annual North American Conference in Chicago in 2011. Dr. Goldhaber-Fiebert's work on this paper was supported, in part, by the National Institute on Aging [K01-AG037593] and through the Center on Advancing Decision Making in Aging (CADMA) and the Center on the Demography and Economics of Health and Aging (CDEHA) at Stanford University. Dr. Bhattacharya and Dr. Garber's work on this paper were also partially supported by the National Institute on Aging [7R37AG036791]. This RCT was registered in the American Economic Association Registry for randomized control trials under Trial number AEARCTR-0000784. Please address correspondence regarding the paper to Dr. Jeremy D. Goldhaber-Fiebert (jeremygf@stanford.edu) 117 Encina Commons, CHP/PCOR, Stanford University, Stanford, CA 94305-6019 USA. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Nudges in Exercise Commitment Contracts: A Randomized Trial
Jay Bhattacharya, Alan M. Garber, and Jeremy D. Goldhaber-Fiebert
NBER Working Paper No. 21406
July 2015
JEL No. D6,I1,I12

ABSTRACT

We consider the welfare consequences of nudges and other behavioral economic devices to encourage exercise habit formation. We analyze a randomized trial of nudged exercise commitment contracts in the context of a time-inconsistent intertemporal utility maximization model of the demand for exercise. The trial follows more than 4,000 people seeking to make exercise commitments. Each person was randomly nudged towards making longer (20 weeks) or shorter (8 weeks) exercise commitment contracts.

Our empirical analysis shows that people who are interested in exercise commitment contracts choose longer contracts when nudged to do so, and are then more likely to meet their pre-stated exercise goals. People are also more likely to enroll in a subsequent commitment contract after the original expires if they receive a nudge for a longer duration initial contract. Our theoretical analysis of the welfare implications of these effects shows conditions under which nudges can reduce utility even when they succeed in the goal of promoting habitual exercise.

Jay Bhattacharya
117 Encina Commons
CHP/PCOR
Stanford University
Stanford, CA 94305-6019
and NBER
jay@stanford.edu

Alan M. Garber
Provost
Harvard University
Massachusetts Hall
Cambridge, MA 02138
and NBER
alan_garber@harvard.edu

Jeremy D. Goldhaber-Fiebert
117 Encina Commons
CHP/PCOR
Stanford University
Stanford, CA 94305-6019
jeremygf@stanford.edu

A data appendix is available at:
http://www.nber.org/data-appendix/w21406
A randomized controlled trials registry entry is available at:
https://www.socialscienceregistry.org/trials/784

1 Introduction

Resolutions to initiate or intensify an exercise program fail nearly as often as they are voiced. But failures to form exercise habits do not generally stem from a lack of desire. Many people wish to increase their physical activity to become healthier, and try – repeatedly – to make exercise a habit, often joining a gym as part of a New Year's resolution. When their resolve fails, it is not simply because they have decided the goal is unimportant. More often, the desire remains, but they have not found a way to convert initial motivation into a sustained habit of physical activity.

One explanation for this mode of failure, advanced in the behavioral economics literature, is that people are time-inconsistent (Laibson, 1997). That is, tomorrow's evaluation of the utility from a stream of future consumption will not match today's evaluation because of an overemphasis on present consumption when tomorrow becomes today (i.e., their pattern discounting is hyperbolic). In practical terms, hyperbolic discounting leads to overweighting of near-term costs and underweighting of delayed benefits. Some people, aware of the consequences of their time inconsistency, will seek constraints that enable them to stick to their resolutions. A commitment contract is one approach, in which the present self offers inducements to future selves to act a certain way (such as exercising more frequently). Just as some parents have favorite children, people pursue commitment contracts because they have favored future selves. An online market has formed around this idea with firms such as stickK.com that broker such contracts.

Just because time-inconsistent people may realize they have a problem keeping to their commitments – and hence demand a pre-commitment device – does not imply that they will inevitably choose a commitment contract that is most effective for forming habits. Behavioral economists, such as Thaler and Sunstein (2008), have proposed the nudge as a solution to related problems. Nudges are small, easily reversible changes to the environment that "gently" push people into a default choice that will obtain unless the person actively switches away from it. Studies have shown, for instance, that nudges can be effective in establishing savings habits. By changing the default to opt-out instead of opt-in, workers were more likely to contribute to retirement savings programs (Thaler, 1994).

Much of the literature on nudges has focused on one-time decisions with day-to-day consequences that become largely invisible.[1] However, many important habits – like exercise – do not have this feature, but instead require substantial ongoing effort to maintain. We study whether nudges can be effective in establishing exercise habits as an example of such a context. To this end, we analyze a randomized trial conducted by stickK.com, and we assess the welfare implications of nudges for time-inconsistent individuals.

[1] Charness and Gneezy (2009) and Acland and Levy (2013) are exceptions to this; they show that short-term cash payments can induce habits in domains requiring ongoing maintenance; sustaining such habits appears challenging.

One definition of a nudge is an intervention that induces a person to voluntarily choose one option over another, where both would produce virtually the same utility. In this view, nudges cannot diminish individual welfare, and may produce positive externalities. The welfare impact of nudges for time-inconsistent individuals is more subtle. A nudged behavior may produce no perceived change in total utility from the point of view of the current self, but induce large changes in behavior – and increases or decreases in total utility – for future selves. Similarly, a nudge that produces a long-term exercise habit may increase welfare for some future selves and may decrease it for others, even though the initial choice to participate in the commitment contract is voluntary.

Our purpose in this paper is to address open questions on the efficacy of nudges to induce exercise habits and improve well-being. Does a relatively short-term commitment to additional exercise – induced by a nudge – lead to habit formation (Stigler and Becker, 1977; Murphy and Becker, 1988)? Do the welfare benefits of the resulting exercise habit (e.g., through improved health and human capital) offset the welfare costs of acquiring the habit for a time-inconsistent individual? Finally, what does it mean to improve the welfare of a time-inconsistent individual whose various selves disagree about how to evaluate any given consumption path?

We begin by asking an empirical question: do nudges change exercise behavior? We describe the experiment we use to test the hypothesis that a nudge to a longer commitment contract will result in more sustained exercise. Given that we answer in the affirmative, we seek to understand the welfare implications of the result using a theoretical framework that explores nudges, time-inconsistency, and welfare. This analysis enables us to consider the welfare consequences of nudges in the context of voluntary commitment contracts for time-inconsistent individuals.

2 A Randomized Nudge for StickK.com

We study whether nudges, in the guise of alternative versions of exercise pre-commitment contracts, promote longer-term exercise habits. In the context of a free, web-based commitment contract site (stickK.com), we conduct an experiment in which we randomly assign users to different suggested lengths of an exercise contract, with an intention to shift each user's chosen contract duration (Goldhaber-Fiebert, 2010). We measure the effect of longer duration nudges on: 1) the pre-commitment contract features that the subject chooses; 2) exercise successfully completed during the contract; and 3) demand for subsequent exercise contracts.

The data we analyze describe users of stickK.com, adults aged 18-75 who made exercise commitment contracts between October 2010 and April 2012. Interested individuals visited the website and constructed contracts. The key features of the contract are its duration – how many weeks the individual commits to exercise – and its frequency - the number of exercise sessions per week (**Figure 1**). Users also selected an amount of money to place at risk of forfeiture for each week in which they failed to exercise according to the contract. Success and failure could be self-reported, but users could also designate a referee to confirm their weekly reports to reduce potential cheating.

In the experiment, we randomly assign the user to an 8, 12, or 20-week default duration for the contract, which can be modified with a simple mouse click.[2] The user also chooses a contracted exercise frequency. We do not randomize the default frequency seen by the user. Next, the user designates the financial penalty for every failed week, which can take a value of zero or any positive number of dollars. Finally, the user can designate a referee to confirm whether she or he has successfully completed each week of exercise.

For each user-initiated contract (signed or not), we collect information on the date, the randomized nudge, the choices made about contract features (duration, frequency, penalties, referee, etc.), whether the contract was signed, and if the contract was signed, the user's weekly exercise performance. Each week that was not a success resulted in a deduction from the user's credit card for the contracted penalty amount. Additionally, we collect baseline information about each user's age, sex, and country (determined by their IP address) along with any prior contracts with stickK.com.

3 The Effect of the Randomized Nudge on Exercise Habits

Overall, 8,809 users initiated exercise contracts, of whom 74% signed a contract. Among those who initiated a contract, 40.9% were male, 33.4 years of age on average, and 68.5% accessed the website from a computer with a US IP address. Few had made any prior exercise commitment contracts with stikK.com (1.7%). The randomization successfully balanced these observable characteristics (**Table 1**).

3.1 Effect of Nudges on Signing an Exercise Commitment Contract and its Characteristics

Longer duration nudges increased the duration of contracts without reducing the likelihood of signing a contract or altering choices of other contract characteristics. The fraction of users who ultimately signed contracts did not differ significantly across different nudges. Among users who signed contracts (**Table 2** and **Figure 2**), longer duration nudges increased the length of chosen contract duration (12.5 weeks for the 8-week nudge; 14.0 weeks for the 12-week nudge; and 18.8 weeks for the 20 week nudge, $p<0.001$). By contrast, the groups did not differ significantly in terms of exercise frequency (~4 times per week), financial stakes, and use of a referee (~31%). Users committed to financial stakes in 22.7% of the signed contracts, with approximately $23 on average at risk for each week that they failed.

[2] For users who initiated more than one contract, randomization was on a per-contract basis.

3.2 Effect of Nudges on Exercise Performance within the Contract Period

Next, we analyze how the randomized duration nudge affected exercise performance during the contract. We restrict our sample for this analysis to users who had signed contracts 26 weeks prior to April 2012. This restriction ensures that we can follow users for up to six months within the contract period (n=3,397).

We first conduct an intent-to-treat (ITT) analysis in which we classify people by their random nudge – regardless of the contract duration actually chosen – and compare outcomes based on this classification. The results from this analysis are shown in **Figure 3**. Users randomized to the 8-week nudge successfully completed 3.2 weeks of exercise [95% confidence interval (CI): 2.9 - 3.5 weeks]. The 12-week and 20-week nudges increased this to 3.4 weeks [95% CI: 2.9 – 4.0 weeks], and 4.3 weeks [95% CI: 3.8 - 4.9 weeks] respectively with only the effect of the 20-week nudge reaching statistical significance at the p=0.05 level. Overall, we find that users randomized to longer duration nudges successfully complete more weeks of contracted exercise.

The ITT analysis is interesting because it shows how nudges affect mean exercise levels. It is also interesting to ask whether people who were successfully nudged into longer contracts exercise more. It is possible, for instance, that people who are "tricked" by the nudge into a longer contract might not respond by actually exercising. To address this possibility, we conduct a treatment-on-the-treated (TOT) analysis, in which we measure how an exogenously assigned longer contract alters the outcome, using the random assignment of the nudge as an instrument for chosen contract length.

To make the results easier to interpret, we divide chosen contract length into three discrete and mutually exclusive groups: $short_i = 1(weeks_i < k_1)$, $long_i = 1(weeks_i \geq k_2)$, and $medium_i = 1(k_1 \leq weeks_i < k_2)$. Here, i indexes over users. We arbitrarily pick cut-off values, k_1 and k_2 to be 11 and 19 weeks, respectively. In the appendix, we report a sensitivity analysis where we vary these cut-offs, but our results are robust to these choices.

The correlation between chosen contract length and exercise behavior does not have a simple causal interpretation because users who choose longer contracts may differ in material and unobserved ways from those who choose shorter ones. Our analytic strategy is to use the randomized nudge as an instrument for chosen duration to account for the endogeneity of contract length as a determinant of exercise behavior. Since the nudge is randomly assigned, it meets the requirements for a valid instrumental variable.

We define an ordered categorical variable to reflect chosen contract length as $duration_i = 0 * short_i + 1 * medium_i + 2 * long_i$. Using this definition, we estimate an ordered logit model, in which we measure the relationship between the likelihood of being in a contract duration category, and a vector indicating which nudge the user was randomly assigned, $nudge_i$:

$$P[duration_i = 0 \mid nudge_i] = \Lambda(\varphi_0 + \varphi_2\, nudge_i)$$

$$P[duration_i = 1 \mid nudge_i] = \Lambda(\varphi_1 + \varphi_2\, nudge_i) - \Lambda(\varphi_0 + \varphi_2\, nudge_i) \qquad (1)$$

$$P[duration_i = 2 \mid nudge_i] = 1 - \Lambda(\varphi_1 + \varphi_2\, nudge_i),$$

Here, φ is a vector of the parameters of the ordered logit model that we estimate using standard maximum likelihood methods, and $\Lambda(.)$ is the inverse logit transformation. Based on this first stage analysis, we calculate the predicted probability of being in each duration category as a function of the assigned nudge: $\widehat{short_\iota}$, $\widehat{medium_\iota}$, $\widehat{long_\iota}$.

In a second stage analysis, we regress the number of weeks of successful exercise on these predicted values:

$$exercise_i = \theta_0 + \theta_m\, \widehat{medium_\iota} + \theta_l\, \widehat{long_\iota} + \varepsilon_i \qquad (2)$$

Here, θ is a vector of the parameters of our regression model, and $\widehat{short_\iota}$ is the reference category. We estimate standard errors from 1000 bootstraps samples from our study population on which we perform both stages.

We find that an exogenously assigned longer contract would increase total successful weeks of exercise (**Figure 4**). Users placed in the shortest category of contracts (<11 weeks) would complete 3.3 weeks of exercise successfully [95% CI: 1.3 – 4.5 weeks]. Those assigned medium length contracts (between 11 and 18 weeks) would complete 2.0 weeks of exercise successfully [95% CI: -0.1 – 4.4 weeks], less than in the short category but with a wide confidence interval. By contrast, we find that users placed into long contracts (>18 weeks) exercise successfully significantly more weeks, 6.0 [95% CI: 4.5 – 7.6 weeks], than those placed in short ones.[3]

3.3 Effect of Nudges on Subsequent Demand for Exercise Commitment

Our motivation in this section is to test the idea that longer initial exercise contracts can cause people to develop exercise habits that last beyond the end of their initial contracts. Combined with the main result of Section 3.2 that longer nudges lead to longer contracts, this would imply that nudges, used in the right context, help some people to develop exercise habits.

[3] This non-monotonic pattern – the minimum in exercise levels occurs for the group exogenously assigned to a medium length contract – can easily be explained by chance. However, even if it were not due to chance, the result would be consistent with the theory of nudges that we develop below. In that theory, nudges resulting in medium length contracts can weaken exercise habits, whereas nudges into long contracts can strengthen them (at least for some types of people). The results are robust to alternative duration category cutoffs k_1 and k_2 (see **Appendix Table A1**).

Unfortunately, we cannot directly address the formation of long-term exercise habits with the data we have, since we do not observe whether a user continues to exercise after the end of the initial contract. We do observe whether a user takes up a subsequent commitment contract with stickK.com after the end of the initial one.[4] Because of this data limitation, we test whether longer duration nudges increase the demand for subsequent stickK.com exercise commitment contracts among the set of people who signed an initial contract. The idea is that users who seek a second exercise contract after the end of their first are likely to be still interested in exercising, while users who do not are less likely.[5]

In the entire sample, approximately 6% of users signed a second contract after their first contract expired. Because in some cases our observation window is shorter than the length of the randomized nudges they were assigned for their first contract, we examine the likelihood of signing a second contract within a fixed period after the first contract ends, restricting our sample to only those individuals whom we observe over that time. We use two follow-up periods: 30 days after the first contract ends (n=4,085) and 90 days after the first contract ends (n=3,448).

Our empirical analysis is analogous to the one we report in Section 3.2. We estimate a logit probability model to measure the effect of the nudge on the probability of signing a second contract by day T, the fixed amount of time since the first contract ended (i.e., 30 days or 90 days in our case). We define $t = 0$ to correspond with the expiry date of the first contract. The only covariates in our regressions are the randomly assigned nudge indicators.

$$P[2^{\text{nd}} \text{ contract by } T \mid nudge_i] = \frac{exp(\tau\, nudge_i)}{1+exp(\tau\, nudge_i)}. \qquad (3)$$

Here, τ is a vector of parameters that we estimate, which includes a regression constant. For this ITT analysis (**Figure 5**), we find that those randomized to the 8-week nudge in their first contract had a 2.8% chance of signing a second contract within 30 days [95% CI: 2.1% – 3.9%] and a 4.3% chance of signing within 90 days [95% CI: 3.3% – 5.6%]. For those randomized to the 12-week nudge in their first contract, these chances were 2.8% [95% CI: 2.1% – 3.8%] and 3.9% [95% CI: 3.0% – 5.2%] respectively. Those randomized to the 20-week nudge in their first contract exercised significantly longer: 4.5% [95% CI:

[4] We also do not observe whether a user initiates a second commitment contract in another manner (e.g., informally with a family member). However, since we are interested in the effect of a randomized nudge on subsequent exercise, this additional assumption does not cause us problems unless nudges of different lengths lead people to change their preferences for stickK.com relative to other forms of commitment contracting. We exclude this possibility here because permitting such *ad hoc* changes in preferences would be consistent with any pattern of exercise and pre-commitment contract demand.

[5] This interpretation works as long as the completion of an exercise contract does not transform a myopic person into someone who keeps to her plans. We exclude this as a possibility for the same reasons we do not permit changes in utility function parameters (please see the previous footnote).

3.5% – 5.8%] and 5.8% [95% CI: 4.5% – 7.4%] for signing a second contract within 30 and 90 days with p<0.03.

For our TOT analysis (**Figure 6**), we employ a 2-stage residual inclusion (2SRI) model of the probability of signing second contract as a function of the duration of the first contract. For this model, as before, we use the nudge as an instrumental variable to predict duration, which is endogenous in the second-stage logit model. The 2SRI estimator is analogous to two-stage least squares, but is consistent even when the second stage is a non-linear function of the covariates. It reduces to two-stage least squares when the second stage is linear (Terza et al., 2008). From this model, we calculate the effect of exogenously moving a user from a short initial contract to a medium or long length contract on the probability of signing a second contract.

We find that an exogenously assigned longer contract would increase enrollment in additional exercise commitment contracts after the first contract ends. The probability that users placed in the shortest duration category for their first contracts (<11 weeks) would sign a subsequent commitment contract within 30 days of the end of their first contract is 3.5% [95% CI: 0.2% – 37.8%] and 5.3% [95% CI: 0.3% – 42.3%] within 90 days. Those assigned medium length first contracts (between 11 and 18 weeks) would sign a subsequent contract 1.3% [95% CI: 0.0% – 8.9%] and 1.4% [95% CI: 0.0% – 5.9%] of the time within 30 and 90 days respectively. By contrast, we find that users placed into long first contracts (>18 weeks) would sign subsequent contracts 9.4% [95% CI: 0.7% – 90.6% weeks] and 11.6% [95% CI: 1.1% – 92.5%] of the time within 30 and 90 days respectively.[6]

4 A Brief Retelling of the Theory of Myopia in the Context of Exercise

Our results thus far show that nudges can increase the chosen length of commitment contracts, exercise done during contracts, and demand for subsequent contracts. Under some assumptions, they also suggest that nudges can increase exercise habits. In the remainder of the paper, we consider whether these changes, induced by the nudge to choose a longer contract, make users better off as a result.

To this end, we employ the theory of time-inconsistent behavior (Laibson, 1997). In our context, the theory also needs to account for the demand for exercise commitment contracts along with the development of exercise habits. After introducing this machinery in this section, in Section 5 we discuss some comparative statics predictions of this model, and in Section 6 we develop a working definition of a nudge so that we can explore how changes in the parameters of the model – exercise prices, discount rates, and degree of myopia – affect the demand for exercise and welfare.

---

[6] These findings were robust to difference in cutoffs used for the instrument construction as described in the previous section on exercise success in the first contract (see **Appendix Tables A2 and A3**).

4.1 Laibson's Model of Hyperbolic Discounting

Laibson's (1997) model conceives of inter-temporal choice as a game played among different "selves," one in each period. This contrasts with the standard model of exponential discounting, in which current and future selves agree about how past and present utilities should be weighted. For the selves in periods $t = 1 \ldots \bar{T} - 1$, utility is given by:

$$U_t = u_t + \delta \sum_{\theta=t+1}^{\bar{T}} \beta^{\theta-t} u_\theta \text{ for } t = 1 \ldots \bar{T}. \tag{4}$$

Each self cares about the discounted welfare of future selves, where $\beta \in [0,1]$ is the rate of time discount. High values of β imply that a self cares a great deal about the welfare of his future selves, even in the distant future, while low values imply the opposite.

Myopia is mediated by the $\delta \in [0,1]$ parameter. If $\delta = 1$, U_t simplifies to the standard exponential model of inter-temporal utility that is common in the literature on rational addiction or human capital. Let $u^*(1) = \left\{u_1^{(1)*}, u_2^{(1)*}, \ldots u_{\bar{T}}^{(1)*}\right\}$ be the sequence of per-period utility values that (subject to budget constraints to be described) maximizes utility from the first period self's perspective. Likewise, let $u^*(2) = \left\{u_2^{(2)*}, \ldots u_{\bar{T}}^{(2)*}\right\}$ be the sequence of per-period utility values that maximizes utility from the second period self's perspective.

A standard result in this setting is that the plan that each self makes for the actions of future selves is maintained when each future self takes control (e.g., Becker and Murphy, 1988; Lucas, Stokey, and Prescott, 1989). There is no real difference between the "selves" since they agree on the optimal plan at all points.

$$u_t^{(1)*} = u_t^{(2)*} \ldots = u_t^{(s)*} \; \forall \, s \leq t \tag{5}$$

However, if $\delta < 1$, the selves may disagree about the optimal plan, and equation (6) no longer necessarily holds. Every current self believes that each future self will care more about selves yet further in the future than they actually end up caring. In our context, the current self will tend to make plans for future exercise that turn out to be overly optimistic. For instance, the period 1 self of someone with $\delta < 1$ might have plans to exercise a lot in period 3, but this plan is foiled after the period 1 self loses control: $u_3^{(1)*} > u_3^{(2)*} > u_3^{(3)*}$.

4.2 Sophisticated and Naïve

There are two types of hyperbolic discounters, those who are naïve about the disagreement among the selves about the future, and those who fully understand that future selves will thwart whatever plans were made earlier once they take control. Both naïve and sophisticated hyperbolic discounters end up exercising less than someone who is not myopic ($\delta = 1$) but whose preferences are otherwise exactly the same (Gruber and Koszegi, 2001).

Yet, given only information about observed exercise levels over time, it is impossible to distinguish between a non-myopic individual ($\delta = 1$) with a low rate of time discounting, β, and someone who is myopic ($\delta < 1$) but with a higher value of β. This is because increases in both β and δ tend to increase the demand for exercise, all else equal.[7]

However, if we observe some additional information about the demand for self-control devices – such as a stickK.com contract – we can identify what sort of preferences someone has. A naïve hyperbolic discounter, failing to recognize his time-inconsistency, does not demand a pre-commitment device. Similarly, a non-myopic individual will have no demand for a stickK.com contract since she correctly realizes that she will stick to her original plans. Only a sophisticated hyperbolic discounter, who realizes that his future selves will ignore his plans, will want a self-control device like a stickK.com contract to discipline his future selves. Since, in this paper, we are analyzing a population of people who have chosen a stickK.com contract, we assume throughout that everyone in our sample is myopic and sophisticated.[8]

5 Exercise Demand by the Time-Inconsistent

In this section, we develop and apply the model of hyperbolic discounting to the problem of exercise habit formation. Our goal is to explore how differences between people in the rate of depreciation of exercise capital and extent of myopia determine optimal exercise paths. To this end, we specify a tractable yet flexible parameterization of a utility function that permits the development of exercise habits, and then conduct a multi-period comparative statics analysis.

5.1 A Tractable Utility Function for Exercise

We specify per-period utility as a quadratic form in (c_t, e_t, k_t):

---

[7] It may be possible to distinguish naïve myopic and non-myopic individuals if we assume an arbitrary functional form for time discounting, such as $f(\theta) = \beta^{\theta}$. However, if we do not make such a functional form assumption, and instead assume only that the time discounting function of the non-myopic person is monotonically declining in time θ, then any optimal exercise path generated by a myopic individual can be replicated by a non-myopic individual with lower rates of time discounting.

[8] It is interesting to consider a case where a naïve hyperbolic discounter is tricked into signing a commitment contract. Absent such a trick, *ex ante* a naïve hyperbolic discounter would have no demand for a pre-commitment contract. *Ex post*, if he were tricked into signing a contract, he would behave in ways that are similar to a sophisticated hyperbolic discounter within the contract. Facing higher costs of not exercising the naïve hyperbolic discounter would exercise more than he otherwise would have. This approach could be appealing, for instance, to paternalistic employers.

$$u_t(c_t, e_t, k_t) = [c_t \quad e_t \quad k_t \quad 1] \begin{bmatrix} \alpha_{cc} & \alpha_{ce} & \alpha_{ck} & \alpha_c \\ \alpha_{ce} & \alpha_{ee} & \alpha_{ek} & \alpha_e \\ \alpha_{ck} & \alpha_{ek} & \alpha_{kk} & \alpha_k \\ \alpha_c & \alpha_e & \alpha_k & 0 \end{bmatrix} \begin{bmatrix} c_t \\ e_t \\ k_t \\ 1 \end{bmatrix} \quad (6)$$

Here, (c_t, e_t, k_t) is the amount of non-exercise consumption, exercise, and exercise-related capital in period t. The user faces a period-budget constraint $c_t + e_t \leq 1$, which assumes that the relative price of exercise to consumption is 1, and that the user has a constant income stream that cannot be transferred across periods.[9]

We model exercise habits with the notion of exercise-related capital, k_t, which reflects past investments in exercise. Like other forms of capital, its ongoing value is maintained or strengthened by continued investment in it:

$$k_t = \gamma k_{t-1} + e_t \text{ for } t = 1 \ldots \bar{T} \quad (7)$$

The parameter, $\gamma \in [0,1]$, reflects the depreciation rate of exercise-related capital. Larger values of γ imply that exercise capital depreciates more slowly. We allow $k_0 \in [0,1]$; in the initial condition, users may enter the model with some exercise capital.

We assume that utility is increasing in c_t, decreasing in e_t , and increasing in k_t. Therefore, we restrict the signs of parameters as follows: $\alpha_c \geq 0$, $\alpha_{cc} \geq 0$, $\alpha_k \geq 0$, $\alpha_{kk} \geq 0$, $\alpha_e \leq 0$, $\alpha_{ee} \leq 0$, and $\alpha_{ek} \geq 0$. The cross-derivative assumptions guarantee that a feasible maximum exists, though the maximum may involve corner solutions. We normalize the α's such that the absolute value of each component of α is less than one.

5.2 Comparative Statics of Myopia and Exercise Capital Depreciation

Our goal in this section is to explore how people with different capacities to build exercise habits (γ) and different levels of myopia (δ) differ in their chosen exercise levels over time. In Section 6 of the paper, we explore how nudges alter optimal exercise paths.

To simplify the exposition and our calculations, we consider a four-period version of the Laibson hyperbolic discounting model (that is, we set $\bar{T} = 4$) with the quadratic preferences that we describe in Section 5.1. Our motivation for this choice is that we want to allow a period prior to making a commitment contract as well as both immediate and longer-term post-contract follow-up as a way to model habit maintenance. Larger values of $\bar{T}$ complicate the analytics without altering the substantive conclusions.

Let $\Pi = \{\gamma, \delta, \beta, \alpha, k_0\}$ designate all the parameters of our model, which include the rate of depreciation of exercise capital (γ), extent of myopia (δ), discount rate (β), preference parameters (α), and initial exercise capital (k_0). Given a particular parameter set, Π, we would like to determine optimal exercise levels, $e^*(\Pi)$, in each period

[9] We normalize c_t and e_t to be between zero and one.

$$e^*(\Pi) = [e_1^*(\Pi)\ e_2^*(\Pi)\ e_3^*(\Pi)\ e_4^*(\Pi)] \quad (8)$$

However, solving for $e^*(\Pi)$ analytically for all Π is challenging because standard methods of dynamic programming do not work with hyperbolically discounted preferences (Gruber and Koszegi, 2001). Because we seek to calculate $e^*(\Pi)$ for a wide range of values of Π, we employ a backwards-induction method, starting from $\bar{T} = 4$, that accounts for the time-inconsistency in preferences. An additional complication, which we address, is that optimal exercise levels may fall to a corner; there are many people whose dislike of exercise leads them never to do it. We use Mathematica to derive the optimal levels of exercise and consumption in each period given Π.[10]

We sample randomly from the parameter space, Π. In particular, we draw 50,000 samples from a uniform distribution over the support of each parameter (though we fix the discount rate at $\beta = 0.97$). For each parameter set, we then solve for $e^*(\Pi)$.

In **Figure 7**, we show how optimal exercise level, $e^*(\Pi)$, changes for different value of δ and γ, integrating over the joint distribution of the remaining parameters in Π. The figure shows how the demand for exercise in each period – $e_1^*(\Pi)$…$e_4^*(\Pi)$ – depends on δ and γ. On average, optimal exercise is lower when people are more myopic (δ is closer to zero) and when people have more difficulty forming exercise habits (γ is closer to zero). Small changes in habit depreciation lead to large changes in optimal exercise levels, especially for exercise in periods 1 and 2. By contrast, even large changes in myopia lead to small changes in optimal exercise levels.

The patterns shown in **Figure 7** arise for two reasons. First, when there is an interior solution for $e^*(\Pi)$, this solution involves more exercise for people with larger values of δ and γ. Second, people with smaller values of δ and γ, are more likely to optimally choose not to exercise at all – especially in the last period when there is no future benefit at that point from an exercise habit.[11] Despite the fact that zero exercise is very common in real-world population, we cannot comment on the relative frequency of Π in the population based on this analysis.

6 Nudges

In this section, we formalize our intuition that small nudges do not change private welfare or utility levels, but may lead to large changes in current and future exercise behavior because of the development of exercise-related capital. In turn, such behavioral changes may have consequences for the welfare of future selves.

6.1 Definition of a Nudge

[10] Please see Appendix 2 for an annotated version of our Mathematica and Stata code.

[11] Please see **Appendix Figure A1** for an illustration of both interior and corner solutions.

A nudge here is an exogenous change in the environment that induces a self to select a longer exercise pre-commitment contract. There are limits to an individual's susceptibility to nudging. In particular, if a nudge would move a self to a path whose utility was sufficiently lower than the non-nudged optimal path, then the self would reject the nudge. Definition: a nudge is perturbation at time t from an un-nudged optimal consumption bundle (c_t^*, e_t^*) to a nudged bundle $\left(c_t^{nudge}, e_t^{nudge}\right)$ such that period t's total discounted utility is reduced by at most an arbitrarily small amount. That is, given an $\varepsilon > 0$, we can conclude that:

$$U_t^* - U_t^{nudge} < \varepsilon \tag{9}$$

Related to our definition of a nudge, we define a *nudgeable range* as the set of exercise and consumption bundles in time t such that utility is within ε of U_t^*. The nudgeable range, of course, will vary with ε, and in general includes a larger set of exercise and consumption bundles for larger values of ε.

Figure 8 illustrates, in a stylized way, our definition of a nudge and its associated nudgeable range. In the figure, the period 1 self chooses e_1^* as an optimal level of exercise, which corresponds to U_1^*. Consider the set of utility outcomes that are within ε of this maximum. The corresponding domain of e_1^{nudge} produces the set of utilities that comprises the nudgeable range.

Figure 8 also shows the effect of nudges in period 1 within the nudgeable range on e_4, the amount of exercise chosen in period 4. Within this set, there is a subset of possible e_1^{nudge}'s that produce a large increase in e_4. In these cases, the exercise related human capital inherited by the period 4 self (k_4) is increased due to the nudge. By definition, the period 1 self chooses more exercise, which ultimately increases k_4. This increase in k_4 in turn leads to greater levels of exercise, e_4, by the fourth period self since $\alpha_{ek} > 0$. A colloquial way describe this is that the early nudge induces an exercise habit that increases the marginal utility of exercise in later periods.

6.2 Nudges and Welfare

As we have seen, a change in the optimal consumption bundle at time 1 due to a nudge can lead to changes in the future consumption path for $\tau = 2 \dots \bar{T}$. Since U_1^* represents the optimal un-nudged utility, the left hand side of equation (9) is always greater than zero; a nudge always reduces the utility of the nudged self, but only by a small amount. However, each future self will re-optimize in response to the downstream effects of the nudge on exercise in period 1 that changes the level of k_τ inherited by the selves, $\tau = 2 \dots \bar{T}$.

A nudge, then, has by definition very little effect on the present self's total utility, but can lead to large changes in the consumption bundles of the future selves and in their total utilities. The model places no restriction on the effect of a nudge on the utility of future selves, which can decrease or increase by amounts greater than ε.

In other words, if a nudge is applied to the current self such that he pre-commits to a longer exercise contract which binds on his future selves, he views this nudged path as equivalent (or nearly so) to the path he would have chosen absent the nudge. However, as exercise is distasteful (the first order effect has $\alpha_e \leq 0$), the selves in subsequent periods may view the nudged path as welfare reducing. Alternatively, if exercise-related human capital accrues sufficiently, selves in subsequent periods may experience increased welfare along the nudged path ($\alpha_{ek} \geq 0$).

6.3 Comparative Statics of Nudges

In this section, we explore how nudges alter long-term exercise outcomes and welfare for people with different preference parameters, Π. That the effect of nudges on these outcomes differ based on Π should not be surprising, since people and selves differ in how they evaluate the development of an exercise habit.

Our strategy is to start by examining the patterns of $e^*(\Pi)$ over our 50,000 simulation samples that cover the support of Π (see Section 5.2). For many of these simulation samples, the optimal exercise levels are fixed at zero in all periods. For a smaller number of samples, there is exercise in the earlier periods, which declines to zero in the latter periods. And finally, for even a smaller number of samples there is an interior solution for optimal exercise in all periods.

To explore the effects of nudges on individuals with these various patterns of optimal exercise, we select at random examples of Π from each pattern. We then systematically nudge exercise up and down from the optimal level chosen at the end of period 1, binding period 2. We then calculate how subsequent period selves reoptimize their chosen exercise levels in response to the nudged exercise in 2. Finally, we calculate the impact of the nudge on total utility from each self's perspective.

Figures 9 and 10 show two patterns that recur in our samples. **Figure 9** considers a person who optimally never forms an exercise habit in the absence of a nudge. In the upper panel of **Figure 9**, the x-axis is the assigned level of exercise at the end of period 1 while the curves of the figure show the levels of exercise chosen in periods 3 and 4 conditional on exercise in period 2. The bright green vertical line shows optimal un-nudged exercise in period 2. The intersection of this line with the period 3 and 4 lines show the levels of exercise in $e^*(\Pi)$ for each period respectively. Similarly, the light blue and gray vertical lines show small and large nudges to the choice made at the end of period 1; the corresponding intersections show the chosen levels of exercise in periods 3 and 4, potentially changed as a consequence of the nudge. In this case, a small nudge in period 1 induces no change in exercise of period 3 or 4 – they continue not to exercise. However, a large nudge causes an increase in both periods; a habit is now formed and sustained.

In the lower panel of **Figure 9**, we consider the welfare implications of the nudged changes to exercise. The setup is similar to the upper panel except that we plot total utility rather than exercise. Each curve in the plot represent total utility from a given self's perspective. In the absence of a nudge, the chosen level of exercise for period 2 corresponds to the highest point on period 2's total utility curve. In the figure, this is represented by the fact that the bright green vertical line intersects the period 2 curve at its maximum. Nudges, both small and large, reduce period 2's welfare. In contrast, the period 3 and 4 selves achieve higher total utility if period 2 can be induced to exercise more. This happens because the later selves are happy to inherit a higher value of k_t (they are happy to have exercised since $\alpha_k > 0$), and because high levels of k_t may even make it worthwhile to continue exercising ($\alpha_{ek} > 0$).

Figure 10 considers a person whose optimal exercise level is positive in the absence of a nudge.[12] In this case, the small nudge pushes the exercise in period 2 below the optimal level while the large nudge still increases period 2 exercise (upper panel). By contrast, a large nudge increases exercise in period 2, but subsequently and perhaps surprisingly reduces exercise to zero in period 3 and 4. The lower panel of **Figure 10** shows that the small nudge to exercise in period 2 decreases welfare for the period 3 and 4 selves. These latter selves attempt to compensate for a loss of utility from a lower value of inherited k_t by exercising more, but ultimately cannot achieve the un-nudged welfare level. A large nudge reduces the need for latter selves to exercise (the period 2 self has already done all the exercising they want) and thus increases the welfare of the period 3 and 4 selves. The large nudge, in effect, destroys an existing habit, but makes the latter selves better off.

7 Conclusions

To study exercise habit formation, we conducted a randomized assessment of the effects of nudging consumers toward longer exercise pre-commitment contracts. We have three main findings. First, longer duration nudges cause people who are interested in exercise commitment contracts to choose longer contracts and to meet their pre-stated exercise goals. They also cause them to take up subsequent exercise commitment contracts after the expiry of the original. These empirical results suggest that nudges in this context show promise as a means to induce exercise habit formation, which is an important public health goal. This finding buttresses evidence from other public health domains (such as smoking cessation) that manipulating pre-commitment contracts can be effective in modifying behavior (Giné, Karlan and Zinman, 2009).

[12] One may ask why a person who exercises in all periods demands a pre-commitment contract as she appears to have no problem with sustaining a habit. However, in this case, myopia implies that she wants her future selves to exercise more than they would otherwise choose.

However, just because nudges can be effective in changing exercise behavior does not necessarily justify their use. This requires a careful analysis of the welfare effects of nudges. As is well known in the literature on time-inconsistency, only sophisticated time-inconsistent individuals (that is, those who can forecast their future myopia) have a positive demand for pre-commitment contracts. For the rest – naïve and time-inconsistent, or time-consistent – nudges do nothing. Thus, the set of people who may potentially benefit from a nudge depends on how many people in population are both sophisticated and time-inconsistent.[13]

Our theoretical results comment on the welfare implications of nudges on people with differing preferences and constraints. Our second main finding is that longer nudges have heterogeneous effects on the welfare of even a single sophisticated time-inconsistent person. Time-inconsistent individuals consist of a set of selves – each in control at a different point in time – who disagree about the optimal exercise path. We find that nudges may help some selves at the expense of others. The selves that are more likely to be helped are those in the periods after the end of the nudged pre-commitment contract. Some reasons for this include: (1) they are happy to have exercised; and (2) the exercise related human capital built up due to the nudge leads them to want to exercise more. Our results may also tie into the literature on prospect theory (Kahneman and Tversky, 1979) in which regret avoidance can be a powerful motivator.

We should note that, while we build our analysis on Laibson's (1997) hyperbolic discounting model, other models of time inconsistency exist in both the economics literature (Frederick et al., 2002) as well as in evolutionary biology. For instance, Stephens et al. (2004) describe a model of impulsiveness in animal populations, and show that impulsiveness may have evolutionary advantages in terms of food consumption and storage. Berns et al. (2007) discuss the brain mechanisms (such as anticipation and self-control) that provide biological basis for non-exponential discounting functions. Bernheim and Rangel (2004) explain time-inconsistent behavior with a dual model of brain function, in which an impulsive hot brain competes for control with a cold brain that cares about the future. In all of these formulations – and unlike the traditional exponential discounting model – there is a potential for disagreement between different selves (or different parts of the brain) in the assessment of the best action in any given period. Hence, nudges or other devices used to change the choice of a self in one period do not automatically give rise to welfare improvements for all selves, as we find in analyzing our model.

Our final main finding is that the effect of longer nudges on welfare can be different for different types of people. For some, a nudge may move someone on to a healthy exercise path that improves welfare. For others, longer nudges may discourage exercise habits, depending on a person's utility, exercise-related capital, rate of time preference, and

[13] This number has not been measured – to our knowledge – and may be differ across domains (exercise, dieting, procrastination avoidance, saving, etc.). We believe this is an important lacuna in the literature that we hope generates interest among empirical behavioral economists.

myopia parameters. In the absence of a nudge, each type of person has an exercise path that they would choose. Suppose a policy were adopted that nudged everyone in the same way (e.g. a 20-week nudge). Because people differ, the nudged contract length will be greater than the un-nudged length for some and less for others. A longer nudge for the latter can destroy their exercise habit, while it can create one for the former. In other words, welfare can be reduced with the creation of a virtuous new habit and increased with its destruction.

We note two important caveats to our analysis. First our paper focuses exclusively on the internal welfare effects of nudges that change exercise among sophisticated hyperbolic discounters. A sedentary lifestyle may also cause external harm by (for instance) leading to chronic disease, raising health care expenditures, and increasing health insurance premiums. We do not consider this source of welfare consequences in this paper because it is worthy of a full-length analysis of its own (see, Bhattacharya and Sood, 2011).

Second, the hyperbolic discounting model of myopia that we use assumes that one's myopia is fixed, as are all the other utility parameters. In particular, the completion of an exercise contract does not transform a myopic person into someone who keeps to her plans. One might imagine a model where these myopia parameters vary over time or in response to some intervention. Stigler and Becker (1977) caution that this approach runs the risk of *ad hoc* theorizing that is not falsifiable. We thus leave this point for future work.

So is a policy that nudges people toward exercise socially beneficial? Based on our analysis, we cannot conclude that is uniformly so. There are people who do not want to be nudged toward exercise, and would be made worse off from having to do so. There are selves within time-inconsistent people who resist exercise, but would be happy that earlier selves have exercised. On the other hand, there are time-inconsistent people for whom a nudge toward more exercise would be welcomed by almost all selves. An exercise-promotion policy that aims toward Pareto improvement should be designed in a flexible way that involves the latter and leaves the former alone. Further consideration should be given to the fact that promoting public health goals may sometimes be at odds with improving social welfare (such as when a nudge destroys a nascent habit). A policy that is guaranteed to help more people than it harms welfare would require more information about each individual's myopia and exercise preferences than any organization typically possesses.

REFERENCES

Acland, D., & Levy, M. (2013). Naivete, projection bias, and habit formation in gym attendance. Working paper series, Social Science Research Network.

Becker, G. S., & Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy*, 675-700.

Bernheim, B.D., & Rangle, A. (2004). Addition and cue-triggered decision processes. *American Economic Review*, 94(5). 1558-1590.

Berns, G.S., Laibson, D., & Loewenstein, G. (2007). Intertemporal choice--toward an integrative framework. *Trends in Cognitive Sciences*, 11(11). 482-488.

Bhattacharya, J., & Lakdawalla, D. (2004). Time-inconsistency and welfare. National Bureau of Economic Research Working Paper #10345.

Bhattacharya, J. & Sood N. (2011). Who Pays for Obesity? *Journal of Economic Perspectives*, 25(1): 139-58.

Charness, G., & Gneezy, U. (2009). Incentives to Exercise. *Econometrica*. 77(3). 909–931

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: a critical review. *Journal of Economic Literature*. 40(2). 351-401.

Goldhaber-Fiebert J.D., Blumenkranz E., Garber A. (2010) Committing to Exercise: Contract Design for Virtuous Habit Formation. National Bureau of Economic Research Working Paper #16624.

Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political Economy*, 223-255.

Gruber, J., & Koszegi, B. (2001). Is addiction" rational"? Theory and evidence. *The Quarterly Journal of Economics*. 116 (4): 1261-1303.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics*, *112*(2), 443.

Lucas, R., Stokey, N., & Prescott, E. (1989). Recursive Methods in Economic Dynamics. Cambridge MA.

Kahneman, D & Tversky, A. (1979) Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2): 263-291.

Stephens D.W., Kerr B., & Fernández-Juricic E. (2004). Impulsiveness without discounting: the ecological rationality hypothesis. *Proceedings of the Royal Society Biological Sciences*. 271(1556):2459-65.

Stigler, G. J., & Becker, G. S. (1977). De gustibus non est disputandum. *American Economic Review*, 76-90.

Terza, J. V., Basu, A., & Rathouz, P. J. (2008). Two-Stage Residual Inclusion Estimation: Addressing Endogeneity in Health Econometric Modeling. *Journal of Health Economics*, 27(3), 531–543.

Thaler, R. H. (1994). Psychology and savings policies. *American Economic Review*, 186-192.

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

Table 1. Individual Characteristics Balanced across Randomization Conditions

Characteristics	8-week nudge	12-week nudge	20-week nudge	P-values for groups differences*
N	2,952	2,981	2,876	
Male (%)	40.0	41.6	41.4	0.416
Age (years)	33.2	33.7	33.4	0.187
USA (by IP address) (%)	67.7	67.9	69.6	0.250
Any prior exercise commitment contracts (%)	1.2	2.0	1.8	0.523

* We tested for between-group differences for binary characteristics (male, USA, prior contracts) using logistic regressions with indicators for the groups. We similarly used a negative binomial regression to test for differences in age in years. No differences were significant at the $p<0.05$ level. Information on gender and location in the USA based on IP address available for 8,709 and 7,981 individuals respectively.

Table 2. Only Contract Duration Differs Significantly across Randomization Conditions among Individuals Signing an Exercise Commitment Contract

	8-week nudge	12-week nudge	20-week nudge	P-values for groups differences*
N	2,157	2,212	2,110	
Chosen contract duration (weeks)	12.5	14.0	18.8	<0.0001
Chosen weekly exercise frequency (days)	4.0	3.9	4.0	0.074
Opted for financial stakes (%)	23.1	22.7	22.2	0.764
Weekly amount for those with stakes ($)	23.80	22.76	22.26	0.865
Contract has a referee (%)	30.1	32.1	31.9	0.296
Contract has supporters (%)	12.2	12.8	12.7	0.885

* We tested for between-group differences for binary characteristics (opting for financial stakes, referee, supporters) using logistic regressions with indicators for the groups. We similarly used a negative binomial regression to test for differences counts (duration in weeks) or else used ordered logistic regressions for counts with only a few categories (frequency). We used linear regressions for stake amounts conditional on having chosen non-zero stakes. The only difference significant at the $p<0.05$ level was that for chosen contract duration, the intent of the intervention. The kernel density plots of chosen contract durations by randomized nudge length are shown in **Figure 2**.

Figure 1. Example of the Exercise Commitment Contract Goal Page – Contract Duration and Exercise Frequency

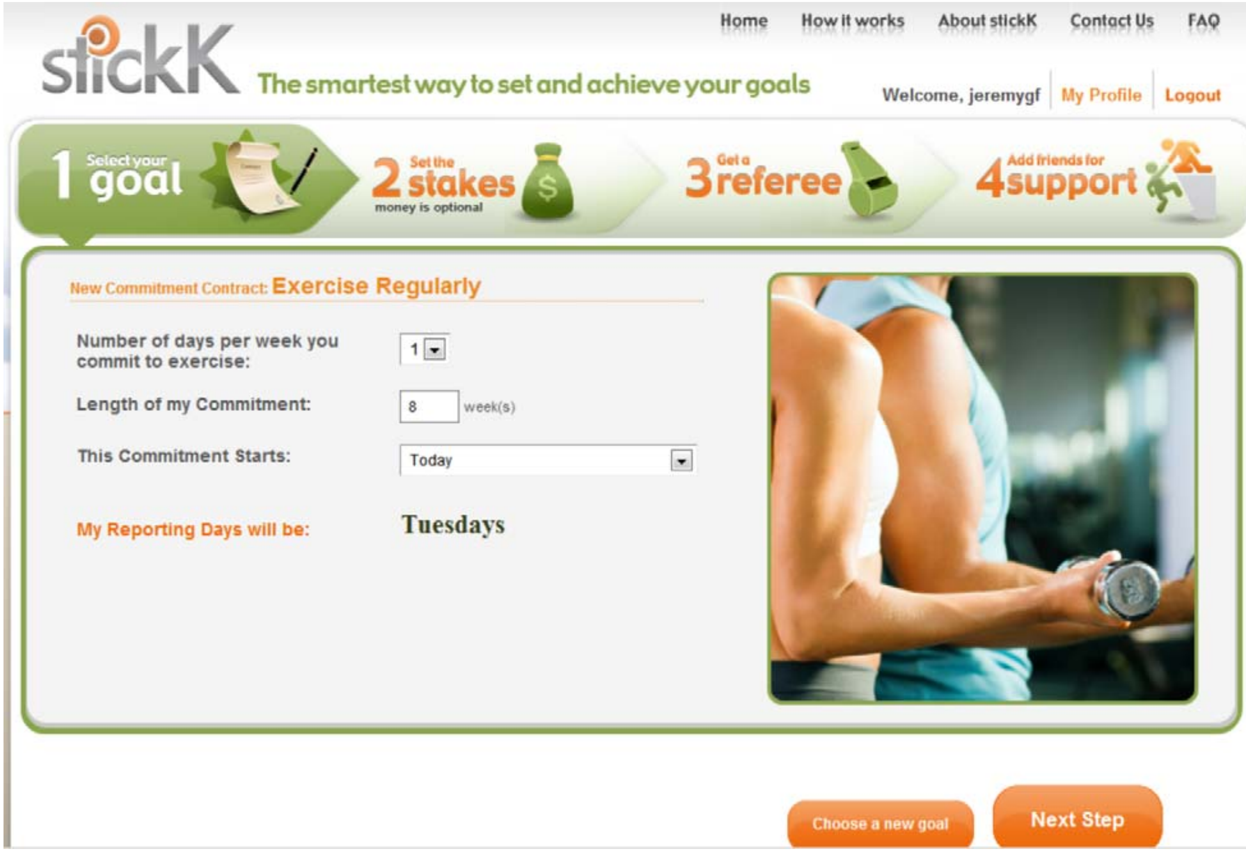

Figure 2. Distributions of chosen contract duration by randomized duration nudge

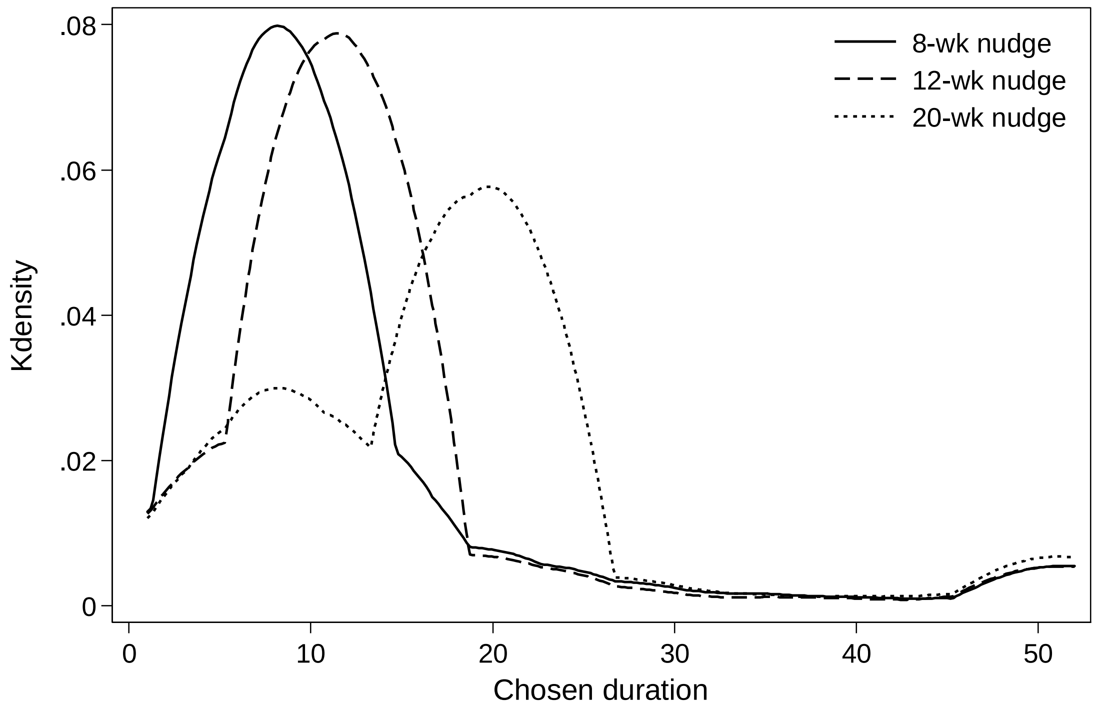

Figure 3. The intention-to-treat (ITT) effect of longer duration nudges on weeks of successful exercise performed during the contract period

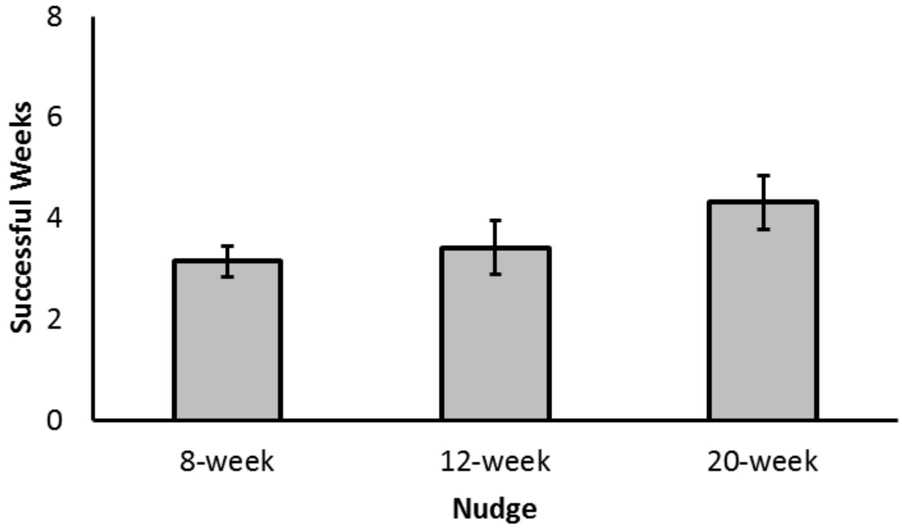

Figure 4. The treatment-on-the-treated (TOT) effect of longer contracts on weeks of successful exercise performed during the contract period

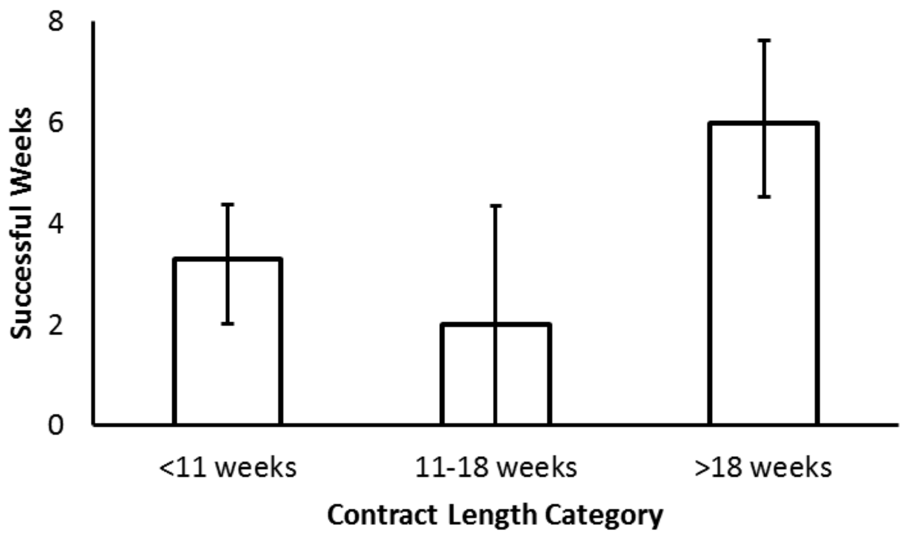

Figure 5. The intention-to-treat (ITT) effect of longer first contract duration nudges on subsequent exercise contract demand within 30 days and 90 days of the end of the first contract

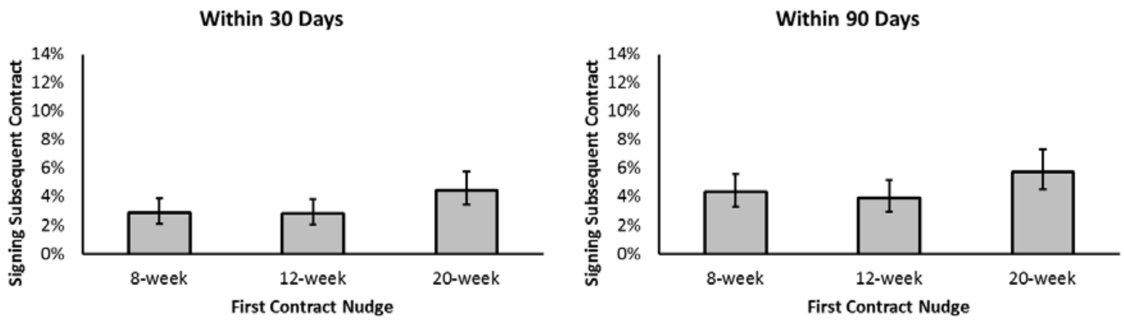

Figure 6. The treatment-on-the-treated (TOT) effect of longer first contracts on subsequent exercise contract demand within 30 days and 90 days of the end of the first contract

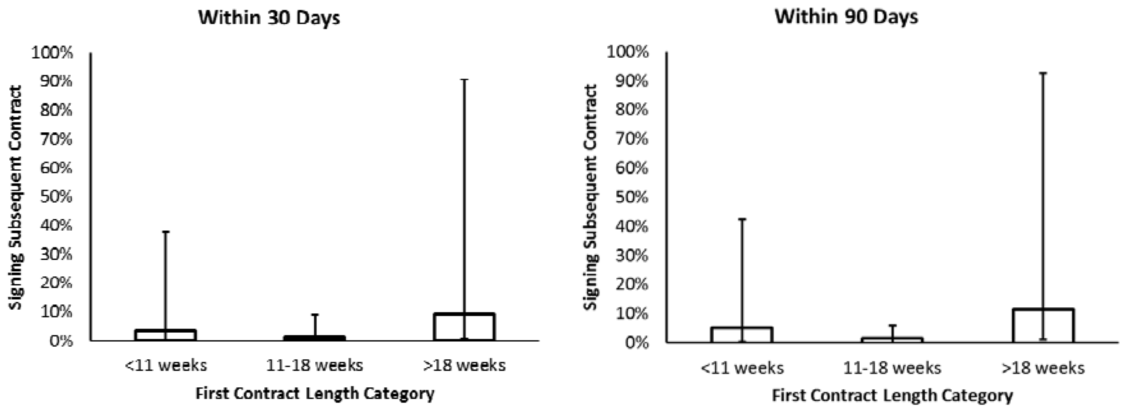

Figure 7. Changes in delta and gamma on optimal exercise (averaged over $f(\beta, \alpha, k_0)$)

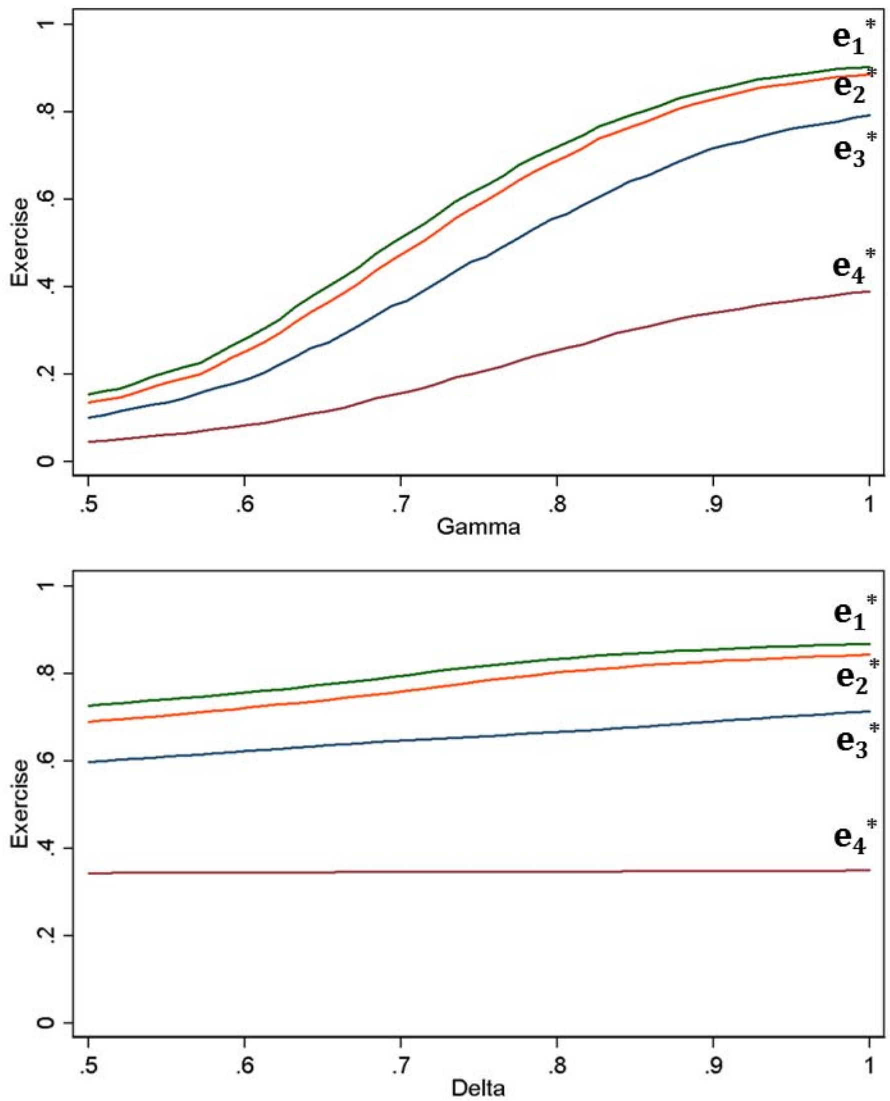

Figure 7 shows how optimal levels of exercise in periods 1 through 4 depend on values of γ (upper panel) and δ (lower panel), averaging over the joint distribution of the other parameters of the utility function (Π) specifically (β, α, and k_0). Individuals who are better able to maintain their exercise-related human capital given their high values of γ and those who are less myopic given their high values of δ have higher optimal values of exercise in all periods on average. Small changes in habit depreciation (γ) lead to large changes in optimal exercise levels, especially for exercise in periods 1 and 2. By contrast, even relatively large changes in myopia lead to relatively small changes in optimal exercise levels.

Figure 8. Illustration of a Nudgeable Range and its relationship to subsequent exercise habit formation

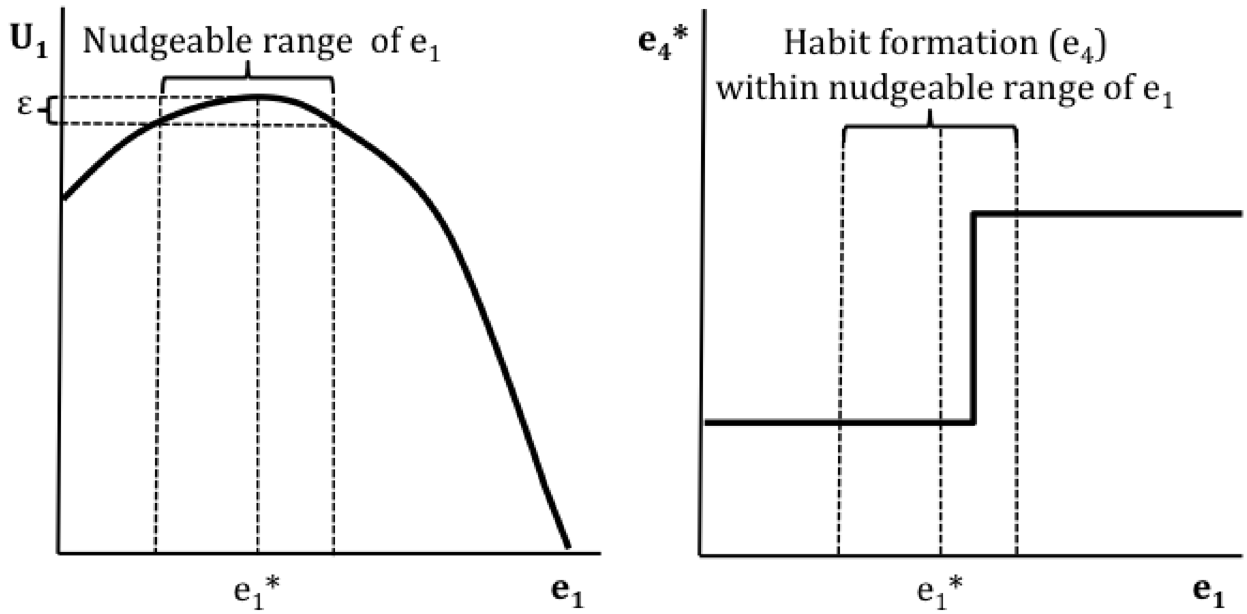

Figure 9. Happy to Have Exercised & Rapidly Decaying Habit

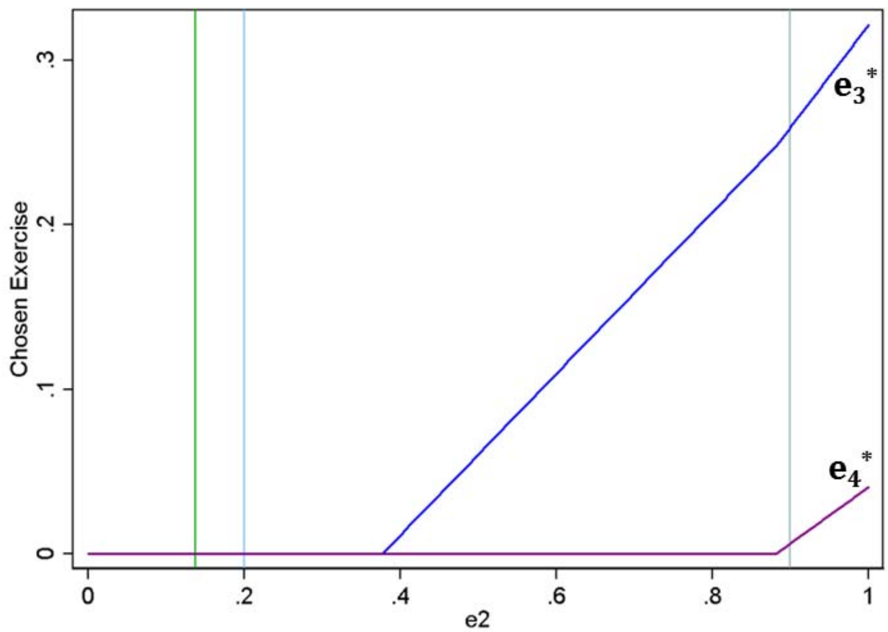

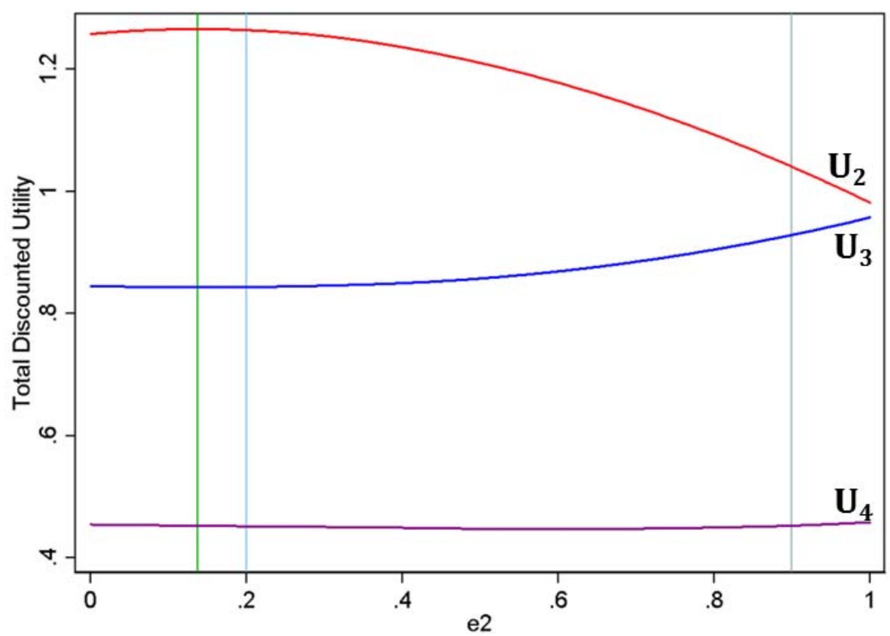

Figure 10. A Longer Nudge Destroys an Exercise Habit – Happy to Have Exercised