



Quantifying the subjective cost of self-control in humans

Candace M. Raio^{a,1} and Paul W. Glimcher^b

^aDepartment of Psychiatry, New York University School of Medicine, New York, NY 10016; and ^bNeuroscience Institute, New York University School of Medicine, New York, NY 10016

Edited by Elke U. Weber, Princeton University, Princeton, NJ, and approved June 17, 2021 (received for review September 5, 2020)

Since Odysseus committed to resisting the Sirens, mechanisms to limit self-control failure have been a central feature of human behavior. Psychologists have long argued that the use of self-control is an effortful process and, more recently, that its failure arises when the cognitive costs of self-control outweigh its perceived benefits. In a similar way, economists have argued that sophisticated choosers can adopt “precommitment strategies” that tie the hands of their future selves in order to reduce these costs. Yet, we still lack an empirical tool to quantify and demonstrate the cost of self-control. Here, we develop and validate an economic decision-making task to quantify the subjective cost of self-control by determining the monetary cost a person is willing to incur in order to eliminate the need for self-control. We find that humans will pay to avoid having to exert self-control in a way that scales with increasing levels of temptation and that these costs appear to be modulated both by motivational incentives and stress exposure. Our psychophysical approach allows us to index moment-to-moment self-control costs at the within-subject level, validating important theoretical work across multiple disciplines and opening avenues of self-control research in healthy and clinical populations.

self-control | decision-making | stress | motivation | precommitment

When Odysseus tied himself to the mast of his ship so he could hear the song of the Sirens without approaching them, he deployed a precommitment mechanism that prevented a self-control failure. When his men were unable to leave the land of the lotus eaters, Homer urges us to see them as having failed in their self-control. However, what does it mean for self-control to fail? This has been a central debate in human behavior for centuries. What has fueled this debate is not a failure to understand what self-control feels like; the subjective experience of resisting temptation is a universal one for humans. What has made self-control so elusive is determining how to convincingly and quantitatively measure it and therefore to understand why it often fails. Whether we are trying to lose weight, quit smoking, avoid drugs, exercise more, drink less, or simply focus on a cognitively demanding task, the question remains: If one truly desires a particular long-term outcome, why is it so difficult to choose in favor of that outcome all of the time?

Emerging theoretical accounts from psychology and economics have attempted to untangle this question using economic models of “cost.” Exercising self-control, these accounts propose, is cognitively costly. From this perspective, “failures” of self-control arise from a rational decision process that weighs the benefits of exercising control against its attendant costs. That is, when the cognitive costs exceed their perceived benefit, individuals should disengage from control processes. These “control costs” are thought to stem from the limited cognitive resources available to support the demands of exercising control. As evidence of these costs, economic theories point to the fact that choosers often adopt precommitment strategies, presumably in an effort to reduce the need for self-control (1–4). Yet, the notion of self-control as “costly” remains controversial in the absence of a robust psychophysical methodology for reliably demonstrating and measuring these costs.

Historically, theoretical accounts that have attempted to explain the puzzling disconnect between what we say we want and what we actually do by pointing to the existence of self-control without providing a platform for its reliable demonstration and quantification. One account, emerging from the psychological literature, points to self-control as a top-down regulatory process that inhibits impulsive action in the service of long-term goals or social norms. Informed by findings from classic delay-of-gratification paradigms [e.g., the “marshmallow task” (5–7); but refer also to ref. 8] and theories of ego depletion [(9–12) but refer also to ref. 13], this account proposes that self-control relies on cognitive resources that are depleted the longer they are used. These theories suggest that the motivational or affective state of a chooser influences the availability or functional integrity of these resources. Fatigued or stressed choosers, for example, are often presumed to have more limited cognitive resources for self-control upon which to draw (10, 14–16). This body of work has dominated psychological conceptions of self-control as a form of “willpower” with impulsive or suboptimal choice emerging from a failure or depletion of control resources. While this work aptly captures the subjectively difficult nature of exercising self-control, it has not provided a reliable method to quantify how much control is needed to successfully resist temptation.

A second account from the neoclassical economic literature examines a host of related choice problems including the failure to save money, over-consumption, and procrastination (refer to refs. 17 and 18 for behavioral examples). These economic models view self-control “failures” as preference reversals (1, 5, 18–21) and have generally eschewed the psychological notion of self-control as a hidden and perhaps unnecessary element. When

Significance

The failure to use self-control is a fundamental problem that humans face in daily life. Recent work suggests that these “failures” might be better understood as a rational decision-making process that weighs the benefits of exercising self-control against its attendant cognitive costs. However, we still know little about how to measure these costs or how they change under different circumstances. Across five independent studies, we find that self-control costs can be measured in humans through monetary willingness to pay to avoid temptation and, further, that these costs appear to be sensitive to incentives, stress exposure, and variability in temptation intensity. Our findings may open avenues of research into computational models of self-control that inform psychological, economic, and health policy research.

Author contributions: C.M.R. and P.W.G. designed research; C.M.R. performed research; C.M.R. analyzed data; and C.M.R. and P.W.G. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

¹To whom correspondence may be addressed. Email: cmr385@nyu.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2018726118/-DCSupplemental>.

Published August 26, 2021.

individuals choose in ways that conflict with explicitly stated goals, these choosers are seen as revealing their true preferences through their observed choices. If this is the case however, why do individuals often choose in ways that conflict with previously stated goals, even choosing in ways that appear inconsistent or irrational (1, 2, 22)? Some behavioral economic work has accounted for this paradox with dual-self models that propose that choosers possess (at least) two sets of preferences that are in active competition (2, 23) and temporal discounting models that include dynamic inconsistencies to drive changing preferences (19, 24, 25, 26).

While these models have provided important ways to quantify decision variables related to self-control, they do not fully explain why individuals are inconsistent in their actual choices. One widely influential resolution to this cross-disciplinary puzzle is to hypothesize that the experience of resisting immediate temptation is effortful and aversive. This inherent disutility implies that self-control imposes a cost on choosers, an idea formalized most recently and elegantly by Gul and Pesendorfer (3, 4) who proposed an axiomatic model of self-control. Gul and Pesendorfer proposed that the presence of temptation in an individual's "menu" of choices will impose a cognitive cost, rendering decisions to reject tempting options more difficult. They hypothesized that if choosers know this, they should prefer choice menus that lack tempting options and might even seek to minimize control costs (and maximize utility) by preemptively eliminating tempting options from their choice menu, a phenomenon referred to as "precommitment" (1, 2, 18). Examples of precommitment strategies that limit control costs might include a dieter who is willing to walk an extra block to avoid walking by a local bakery or a gambler who drives an extra hour to avoid passing near casinos. Gul and Pesendorfer argued that preferences for precommitment reveal choosers' subjective cost of exercising control, pointing to a novel decision variable through which these costs can be measured.

What has limited the impact of this set of hypotheses in an empirical sense, however, is the absence of quantitative data to support it. Is there direct quantitative empirical evidence that self-control is costly, as so many have proposed? Are those costs stationary over time? Are these costs influenced by motivational or affective state as the psychological literature has proposed? Do these costs scale with varying levels of temptation? Despite a number of real-world observations of precommitment (18), we still lack an empirical psychophysical tool for answering these questions.

Here, we utilized a psychophysical approach to test the hypothesis that exercising self-control is cognitively costly and that these costs can be measured using a precommitment mechanism. While we acknowledge that there are undoubtedly many components that feed into an overall subjective cost of self-control, our goal here was to simply measure an aggregate of these costs. To do this, we developed an economic decision-making task that measures how much participants are willing to pay to adopt a precommitment device that removes temptation from their choice environment. We refer to the maximum dollar value participants will pay to remove temptation as their subjective "cost of self-control" and show that these costs respond rationally to incentives, scale with increasing levels of temptation, and predict rates of self-control failure. We further test the hypothesis that these costs are modulated by affective state, finding that stress exposure appears to increase the cost of self-control. Finally, we test the hypothesis that self-control costs grow with the ongoing exertion of self-control but find no empirical support for this hypothesis. These data identify a psychophysical approach for the measurement of self-control costs and may open avenues of research into computational models of self-control that may inform psychological, economic, clinical, and health policy research.

Results

In our experiments, healthy, hungry dieters first provided health, taste, and temptation ratings for a series of food items, allowing

us to identify high- and low-tempting foods for each individual. Participants initially reported the most they would be willing to pay (from a \$10 monetary endowment) to avoid having the high-tempting food placed immediately in front of them for a 30-min period (Fig. 1). With a probability of 2%, this reported "bid" was entered into a standard economic auction procedure (27) (*Methods*) that incentivizes participants to report their true subjective value—in this case, the value of eliminating exposure to temptation (precommitment). If they won this auction, the high-tempting food was replaced with a low-tempting food for 30 min. If they lost the auction, the high-tempting food remained in the room with them for 30 min. The price participants were willing to pay provided a within-subject estimate of the cost of self-control; that is, it reflected the maximal dollar value they were willing to pay to avoid exercising control.

The first bidding trial was made without the high-tempting food in the room to capture each participants' prospective estimate of how costly self-control exertion would be (before exposure to temptation). If the initial trial was not realized (as occurred on most trials), the high-tempting food was brought into the room. Participants were then prompted to report, periodically during 30-min of exposure, how much they were willing to pay to replace the high-tempting food with the low-tempting food for the next 30 min. As with the initial bid, these subsequent bidding trials had a 2% chance of going to auction, which would bring the experiment to a premature close. Bidding trials were collected every few minutes (*Methods*) for the 30-min exposure period. If the 30-min exposure period elapsed without any bid being realized, the subjects remained in the room with the high-tempting food for a final 30 min. This allowed us to track 1) how self-control costs changed over time as participants were continuously exposed to temptation and 2) whether self-control ever failed.

The only observable behavior of interest during the final 30-min interval of the experiment was whether the participant consumed the food. This realization phase of the experiment ensured that the task was incentive compatible, meaning participants' choices allowed them to avoid real temptations and any negative outcomes associated with those temptations. The 2% chance of each bid being realized ensured participants knew that what they bid on the current trial was important, since it could determine whether they were required to spend the next 30 min alone in the experiment room with a tempting food reward, the consumption of which did not align with their stated goals.

Study 1: Self-Control Imposes Costs as Revealed by Willingness to Pay for Precommitment. Our primary question of interest was whether the presence of a tempting food that participants want to avoid consuming does in fact impose a cost on choosers. If so, participants should be willing to pay to remove temptation and eliminate the associated control cost. In accord with predictions from these economic models, we found that participants ($n = 32$) were willing to pay a maximum of (on average) 15% of their \$10 endowment (or $\$1.57 \pm 1.78$ SD) to adopt a precommitment device to avoid temptation. This provides a direct scalar measurement of their subjective cost of resisting temptation (Fig. 2A). Participants were not only willing to pay to prospectively to avoid temptation (mean of first bid trial, pre-exposure = $\$1.47 \pm 1.59$ SD), but they continued to pay throughout the task, providing a continuous measurement of the underlying costs of resisting temptation with longer exposure to the food (mean of subsequent bid trials, post-exposure = $\$1.58 \pm 1.82$ SD).

Our initial experiment demonstrates that self-control imposes a subjective cost on choosers that can be measured monetarily. However, the process of deploying self-control in the presence of temptation has been proposed to change over time. Specifically, continued exposure to a tempting good is often thought to increase the difficulty of continuing to exert self-control. We next examined the dynamics of how bids changed both before and

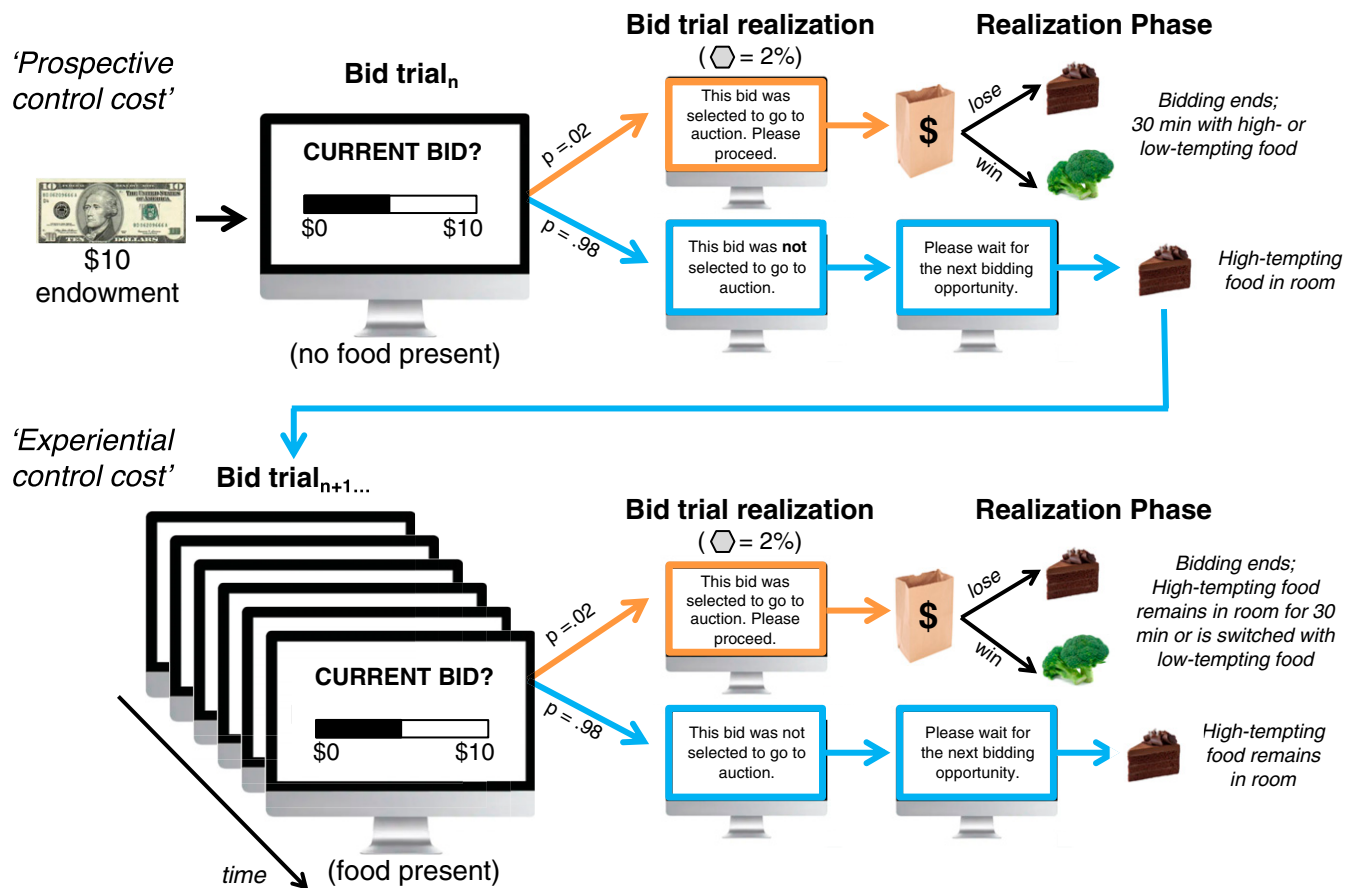


Fig. 1. Illustration of the self-control decision task. Participants reported their willingness to pay to avoid a tempting food reward both before the food was present (*Top*) and, periodically over a 30-min period, with added direct exposure to the food (*Bottom*).

after exposure to temptation (Fig. 2A). We found that average pre- versus post-exposure bids did not significantly differ [paired samples Student's *t* test: $t_{(31)} = -0.533, P = 0.598, d = 0.099$, two-tailed], suggesting that our dieters appear to be accurate in their prospective cost estimates. A repeated-measures ANOVA assessing post-exposure bids as a function of time indicated that on average, there were no significant linear trends in these costs despite ongoing exposure to temptation [$F_{(4,122)} = 0.722, P = 0.576$; Greenhouse-Geisser correction factor, $\epsilon = 0.30; \eta_p^2 = 0.023$]. Thus, we found no evidence that on average, ongoing exposure to temptation increased control costs over the time interval used in our task. We note, however, that individual variability exists in our data set, such that

some participants' bids increased systematically, while for others, they decreased. Finally, we note that 22% of our subjects consumed the tempting food during the ensuing 30-min exposure period ("realization phase"; Fig. 2B).

Study 2: Motivational Incentives Modulate Willingness to Pay for Precommitment. Next, we sought to both replicate this finding and further assess how motivational incentives to sustain goal-directed behavior might affect these costs. In this second experiment, we increased the cost of self-control failure and examined whether participants were then willing to pay more for precommitment to reduce these failures. To do this, we repeated our study

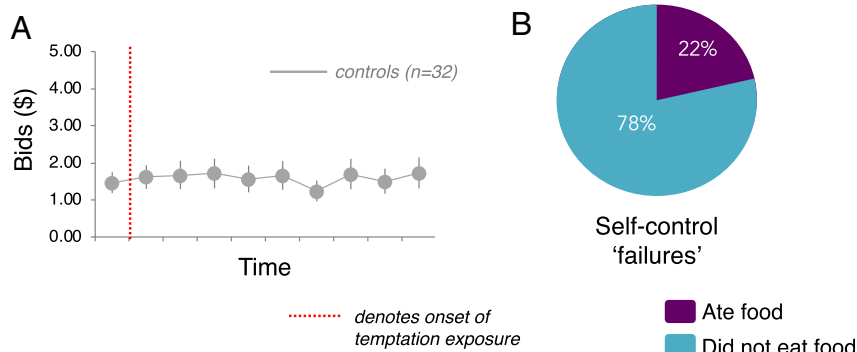


Fig. 2. Study 1. (A) Bids over time for control group (bids 1–10). (B) Proportion of subjects in Study 1 that consumed the tempting food during the study. Error bars denote SEM.

in a second cohort of dieters but instructed participants that they would lose a \$15 bonus at the end of the study if they consumed the tempting food at any point. We hypothesized that by increasing the cost of self-control failure, the value of a precommitment strategy that restricts temptation should also be higher.

In total, 34 new dieters completed our self-control measurement task with the addition of this second monetary incentive. Dieters again showed a reliable and consistent willingness to pay to avoid temptation (mean bid = \$2.85 ± 2.70 SD). Consistent with our hypothesis, the addition of the \$15 cost for eating the tempting food led to a higher willingness to pay for precommitment (Fig. 3A). Combining the data across experiments 1 and 2, a repeated-measures ANOVA with a within-subject factor of time (bids 1 through 10) and a between-subject factor of group (no incentive, incentive) revealed a main effect of group [$F_{(1,64)} = 4.95, P = 0.03, \eta_p^2 = 0.07$] but no effect of time ($P = 0.73$) or time × group interaction ($P = 0.45$). This difference at the group level suggests that participants were willing to spend more money to sustain goal-directed behavior when the costs of not adhering to this goal were higher. We note also that under these conditions, none of the subjects who faced the tempting good consumed it (Fig. 3B).

Study 3: Acute Stress Increases the Cost of Self-Control. Given the tightly coupled relationship between self-control failure and the experience of negative emotional states such as stress, we next examined how exposure to an acute stressor would influence participants' self-control costs. Specifically, we tested the widely held hypothesis that stress makes self-control more "costly." To elicit subjective and neurophysiological stress responses, we recruited a new cohort of dieters ($n = 31$) that underwent the Cold-Pressor Task (CPT) (28–30) prior to the self-control choice task. The CPT is widely used in laboratory settings to reliably induce mild-to-moderate levels of physiological stress and requires participants to submerge their forearms in ice-water continuously for 3 min (Methods). Confirming the efficacy of our stress induction procedures, participants in the CPT condition showed elevated concentrations of salivary cortisol, the primary neuroendocrine marker of Hypothalamic–Pituitary–Adrenal (HPA) axis activity (SI Appendix, Fig. S1A).

We assessed whether our stress manipulation influenced the cost of self-control both prior to temptation exposure—when participants were prospectively estimating these costs—and after the high-tempting food was introduced. Fig. 4A depicts aggregate trial-by-trial bidding behavior for participants in the stress condition. A repeated-measures ANOVA across all studies revealed a main effect of group [$F_{(2,94)} = 4.48, P = 0.01; \eta_p^2 = 0.087$], but no effect of time ($P = 0.45$), or time × group interaction ($P = 0.64$).

Follow-up Student's t tests confirmed that stressed participants reported a higher willingness to pay overall (mean bid = \$3.38 ± 3.04 SD) relative to non-stressed controls [independent-sample Student's t test: $t_{(61)} = -2.88, P = 0.005, d = 0.72$, two-tailed], suggesting that our experimentally induced state of stress was reflected in the valuation of precommitment to restrict temptation. These increases in the stress group were observed during prospective bids [pre-exposure, Bid₁: $t_{(61)} = -3.71, P = 0.0004, d = 0.93$, two-tailed] and persisted across subsequent post-exposure trials [mean of post-exposure bids: $t_{(61)} = -2.77, P = 0.007, d = 0.67$, two-tailed]. Thus, exposure to acute stress appears to have more than doubled (at a between-subjects level) the average subjective cost of self-control. We note that similar to Study 1, 23% of our subjects consumed the tempting good (Fig. 4B).

Study 4: The Effects of Stress and Incentives on Self-Control Costs. To assess whether stress exposure increased self-control costs above and beyond that which we observed when motivational incentives were introduced, we conducted an additional study in an independent cohort of dieters. A total of 31 new dieters completed the self-control decision task with the \$15 penalty for eating the tempting food after undergoing the stress manipulation. We once again observed a reliable willingness to pay to avoid temptation (mean bid = \$2.74 ± 2.20 SD) that was elevated relative to the Study 1 controls (Fig. 5 and SI Appendix, Fig. S2; $t_{(61)} = -2.34, P = 0.023$, two-tailed). To assess whether stress exposure changed bids relative to non-stressed participants who experienced the same incentive structure, we conducted a time × group repeated-measures ANOVA. However, this analysis revealed no main effect of time ($P = 0.11$) or group ($P = 0.86$), nor time × group interaction ($P = 0.65$), thus revealing no significant differences in bidding behavior (Fig. 5). We reasoned that a failure to observe group differences as a function of stress exposure could be due to the fact that the stress manipulation in this particular cohort did not effectively elicit cortisol responses. Indeed, an examination of cortisol concentrations revealed that cortisol levels in this cohort did not differ between groups at any time point (SI Appendix, Fig. S1B). Thus, it appears that while motivational incentives led to an increase in bids relative to control participants (replicating our effects from Study 2), we did not observe differences between the two incentive groups as a function of stress exposure, perhaps due to the failure of the stress procedure.

Secondary Analysis 1: Self-Control "Failures" Were Associated with a Higher Willingness to Pay to Avoid Control. If we assume that subjects do face costs for exercising self-control, then we should expect to see subjects experience self-control failures on occasion. Further, we might expect to find that subjects willing to pay more

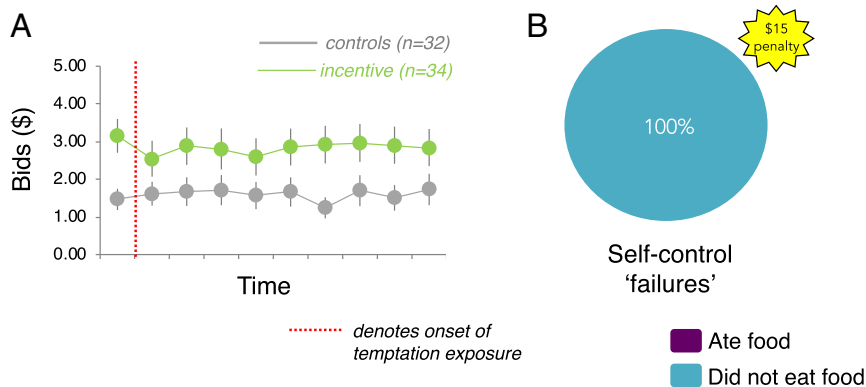


Fig. 3. Study 2. (A) Bids to avoid exposure to the tempting food over time in participants, for which a \$15 monetary loss was imposed for consuming the food (depicted in green; Incentive group) and for those where no monetary loss was imposed (depicted in gray; Control group). (B) Proportion of subjects in Study 2 that consumed the tempting food during the study.

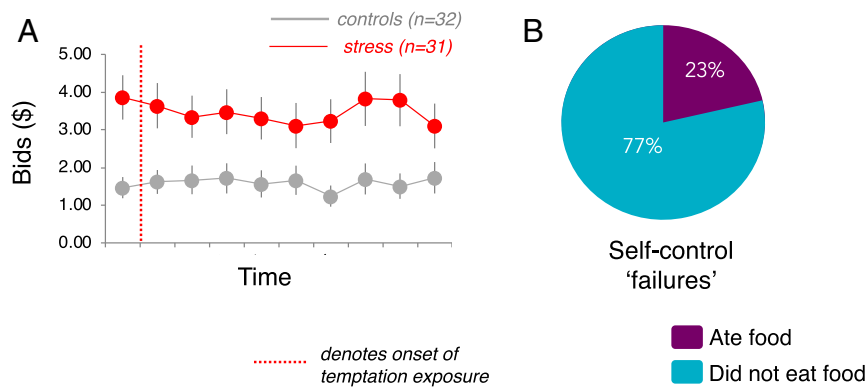


Fig. 4. Study 3. (A) Bids to avoid exposure to the tempting food over time for participants that underwent a physiological stress manipulation (depicted in red; Stress group) and for non-stressed participants (depicted in gray; Control group). (B) Proportion of subjects that consumed the tempting food during the study.

for precommitment experience higher self-control costs and thus might be expected to fail in their self-control more often than subjects who report lower monetary costs for self-control. To test these hypotheses, we examined each subject's behavior after the bidding phase of the experiment was complete—during the final 30-min phase of the experiment (*Methods*). During this phase, if no bids had been realized (which was usually the case given the low probability any bid was realized), participants were required to remain in the room for the final 30 min of the experiment with the tempting food. No further bids were collected during this phase. During this period, we simply observed whether or not each participant consumed the tempting food (a self-control “failure”).

The proportion of participants that consumed the tempting food are presented alongside bidding behavior (Fig. 2B–5B). We focused our comparison on the study cohorts that shared similar penalty structure in order to control for the increased cost of self-control failure. In Study 1 and 3, in which no monetary penalty was imposed for eating the tempting food, 22% and 23% of participants consumed the food, respectively (Figs. 2B and 4B). Given that consumption rates did not differ between these two groups, bidding data were collapsed across these studies to examine how bids differed in dieters that ate the food versus those that did not. Those participants who ate the food were willing to pay significantly more to avoid temptation relative to participants who did not eat the food [independent samples Student's *t* test: $t_{(61)} = 2.81$, $P = 0.006$; two-tailed; Fig. 6]. In Study 2 and 4, in which we

imposed a monetary penalty for self-control failure, no participants consumed the food (Figs. 3B and 5B); thus, a comparable analysis to that above was not possible. We note that the above differences between “eaters” and “non-eaters” remain significant even when including all (“non-eater”) participants from Study 3 [$t_{(95)} = 2.41$, $P = 0.02$] and, additionally, Study 4 [$t_{(126)} = 2.43$, $P = 0.02$].

Secondary Analysis 2: Individual Difference Measures Associated with Control Costs.

We next examined how individual differences across our entire sample related to self-control costs. Given that control costs were higher in stressed participants (Study 3), we first sought to characterize how subjectively perceived stress related to willingness to pay to avoid self-control. To do this, we correlated mean bids and self-reported stress levels directly before the choice task across participants from all four studies ($n = 128$). Perceived stress was positively correlated with average bidding behavior (Spearman's rho: $r = 0.22$, $P = 0.012$; Fig. 7A), suggesting that, across all participants, subjective stress state was related to individuals' willingness to pay to avoid temptation.

Given our evidence that self-control is explicitly costly to choosers, we reasoned that more experience (or success) avoiding temptation might relate to an individual's self-control costs. To explore this question, we conducted a correlation analysis between participants' length of diet and their average willingness to pay to avoid exercising self-control. This analysis revealed a significant negative correlation between mean bids and diet length (Spearman's rho: $r = -0.23$, $P = 0.01$; Fig. 7B). Thus, participants on a diet for a longer length of time

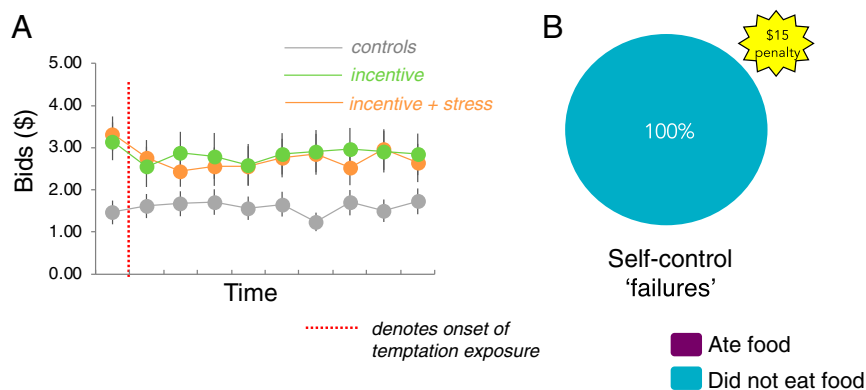



Fig. 5. Study 4. (A) Bids to avoid exposure to the tempting food over time in stressed participants for which a \$15 monetary loss was imposed for consuming the food (depicted in orange; Incentive+Stress group) and for those with the same penalty imposed but no stress induction (depicted in green; Incentive group; Study 2). The control group from Study 1 (no incentive or stress) is depicted in gray for reference. (B) Proportion of subjects from Study 4 that consumed the tempting food during the study.

No incentive groups
(Study 1 & 3)



■ Eat food
■ Did not eat food




Fig. 6. Mean bid for subjects who demonstrated self-control “failures” (23%, depicted in purple) and those who did not (77%; depicted in blue) collapsed across Study 1 (control group) and 3 (stress group). Those participants who consumed the tempting food during the study revealed a higher willingness to pay to avoid control. No participants from Study 2 or 4 (\$15 penalty groups) consumed the tempting food. Error bars denote SEM.

tended to pay less to avoid temptation. We note that this does not reveal whether only those with idiosyncratically lower self-control costs succeed at remaining on a diet, whether self-control costs decline as one’s diet progresses, or both. However, our method applied longitudinally could be used to answer such questions in future work.

Study 5: Willingness to Pay to Avoid Control Scales with Temptation Level. If a participants’ willingness to pay to avoid temptation does in fact reflect the cost of self-control, then we would expect these costs to scale with varying levels of temptation (i.e., when facing a more highly tempting good, a subject should have to exert more self-control than when facing a less tempting good). To test this, we conducted a final study in an independent cohort of healthy, hungry dieters. In this study, participants again rated a series of snack foods on how healthy, tasty, and tempting they were, which allowed us to identify a low-, medium-, and high-tempting food for each individual. On each trial, participants reported their willingness to pay (from \$0 to \$10, from a \$10 endowment) to avoid each of the three food items for varying amounts of time (1 to 60 min; *Methods*). Unlike Studies 1 through 4, all bids were reported prospectively (without any food present), and one trial was randomly selected at the end of the session to be realized.

Fig. 8 depicts average bids for each time point for each level of temptation. A temptation level (low, medium, or high) \times time (1 to 60 min) RM-ANOVA revealed a significant main effect of temptation level [$F_{(2,30)} = 33.06, P < 0.0001, \eta_p^2 = 0.69$] and time

with food [$F_{(9,135)} = 29.12, P < 0.0001, \eta_p^2 = 0.67$] as well as a temptation \times time interaction [$F_{(18,270)} = 5.75, P < 0.0001, \eta_p^2 = 0.27$]. Bids differed significantly for foods with low (mean bid = $\$1.16 \pm 0.36$ SD), medium (mean bid = $\$2.99 \pm 1.02$ SD), and high (mean bid = $\$4.96 \pm 1.30$ SD) temptation levels. Further, bids scaled with each temptation level across increasing amounts of time (Fig. 8). Planned contrasts demonstrated a significant linear increase in bids with higher temptation level [$F_{(1, 15)} = 43.95, P < 0.0001$] and, separately, with increased time with food [$F_{(1, 15)} = 44.81, P < 0.0001$].

Discussion

A universal paradox in human behavior is our tendency to set difficult long-term goals but then to make choices that appear to contradict or undermine those goals. Using an economic decision-making paradigm, coupled with a sample of dieters avoiding tempting food rewards as a model of self-control more broadly, we found that people will pay to avoid temptation, quantitatively revealing their subjective cost of control under a variety of circumstances. We found that these costs are modulated by incentives (i.e., shows cost sensitivity; Study 2/4) and—tentatively—by acute stress exposure (Study 3), consistent with the notion that motivational and affective state modulate one’s willingness to exercise self-control. In a final study designed to test how self-control costs scale with increasing levels of temptation (Study 5), we showed that longer exposure to a tempting good imposes higher self-control

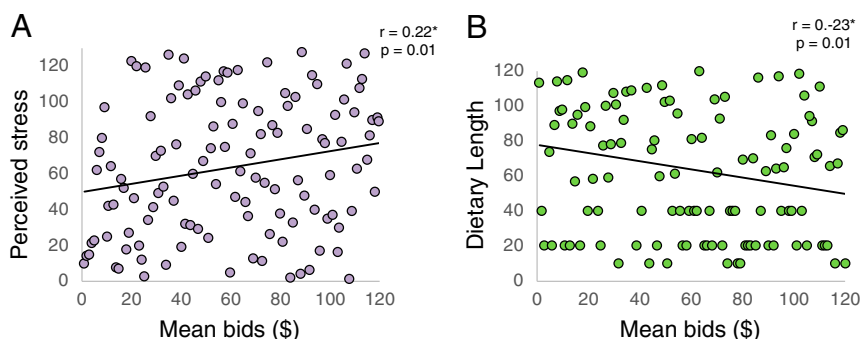


Fig. 7. Individual difference correlations. (A) Perceived stress was positively correlated with average bidding behavior across participants. (B) Length of diet was negatively correlated with average bidding behavior across participants. *denotes Spearman ranked correlation.

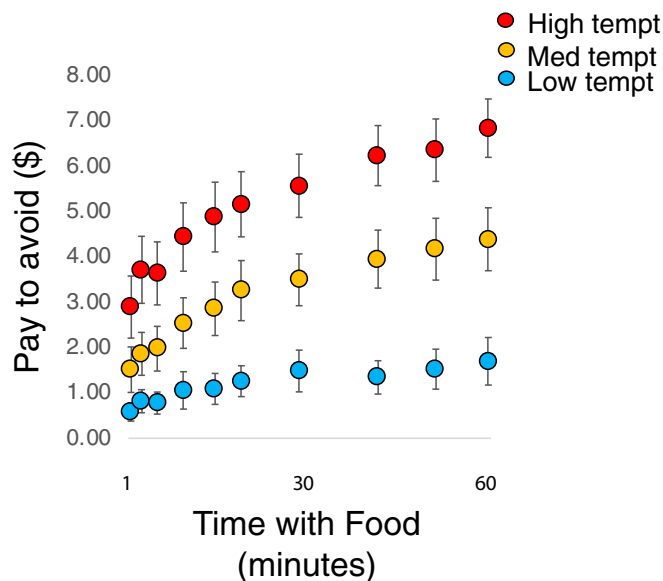


Fig. 8. Willingness to pay to avoid foods that varied in temptation level and amount of time required to spend with food. Bids scaled with increasing time with and temptation level of the food. Error bars depict SEM.

costs (i.e., self-control costs obey monotonicity) and more tempting goods impose higher control costs at the within-subject level than do lower tempting goods (i.e., self-control costs order with temptation level).

Decades of psychological research have revealed that the act of engaging in self-control is subjectively effortful and aversive. Emerging work in the cognitive control literature has proposed that this experience stems from the cognitive costs imposed by deploying control, an account consistent with more recent economic theories of self-control (3, 4) that propose preferences for precommitment reveal an inherent psychological cost to resisting temptation. These converging lines of work provide a clear and testable hypothesis: Self-control failures may be conceptualized as a rational decision that emerges when the costs of exercising control exceed the relative perceived benefits (13, 31, 32). However, without a psychophysical method to precisely quantify these costs, researchers often must infer whether and how much self-control a chooser requires to make goal-consistent decisions or to delay gratification. Our findings unite psychological and economic theories of self-control and provide empirical evidence that self-control is explicitly costly to humans and that these costs can be quantified by measuring the value of precommitment to restrict temptation.

Gul and Pesendorfer's theory is consistent with a growing body of psychological and neuroscience research that suggests people view cognitive demand as intrinsically costly and tend to avoid utilizing cognitive resources if possible (13, 32–34). These costs are thought to stem from limitations in the cognitive resources available to support the demands of control, suggesting self-control arises from evaluations of how valuable expenditures of control are perceived to be relative to how costly. This is a notable departure from classic psychological self-control models that view cognitive resources as depletable (i.e., both limited and diminished with use), arguing instead that such resources are finite and reallocated dynamically depending on the perceived costs and benefits (13, 35). Why some classes of cognitive control may be aversive or costly remains unclear. However, the approach presented here offers a metric for how aversive or costly individuals find the exertion of self-control to be on a moment-to-moment basis.

Our task employed two important features to probe the nature of self-control. First, our task measured momentary willingness

to pay to avoid control prospectively (prior to food exposure) and again after participants encountered the tempting food. One possible explanation for why self-control appears to fail so often is that individuals may poorly estimate how costly self-control will be. Our data tends to lean against this conclusion, at least in this particular choice setting. We observed no significant difference in measured control costs before versus after food exposure. This suggests that our participants had an accurate prospective awareness of the self-control costs they would later face. A second important feature of our task was that unlike existing self-control decision paradigms, our bidding measurements were collected continuously over time, allowing us to track how these costs change with continued exposure to temptation. Some existing work (9, 12) suggests that self-control becomes more difficult as it is continuously exerted. If this were true for our participants, we would expect their self-control costs to increase as the experiment progressed. Interestingly, we instead observed that participants' bids to avoid having to exert control were quite stable over time. The fact that participants' control costs did not increase over time is consistent with a growing body of work suggesting that performance reductions are not a mandatory feature of engaging in control (13, 32, 34). Our findings align with value-based frameworks that argue individuals need not experience decrements in control performance as long as the perceived benefit of deploying control continues to outweigh the cost. We note, however, that our data on this point is relevant only to the intervals of time tested here (i.e., an hour in total duration).

The stability in our participants' willingness to pay to avoid exerting control may also reflect the lack of temporal uncertainty inherent in our task. In a recent line of work, McGuire and Kable (8, 35) demonstrated that behavior in a range of delay-of-gratification tasks (including the "Marshmallow task") (5) might be explained by the underlying predictions participants have regarding when a delayed outcome will arrive. They suggested that one major reason individuals appear to "succumb" to temptation is that under temporal uncertainty, individuals may rationally conclude that the delayed outcome may no longer be worth waiting for. In our study, participants were fully informed regarding the temporal structure of the task and were at all times aware that self-control would be engaged for only a finite period of time. This feature of our task may explain the stability observed in participants' reported bids over time. An open question for future research is whether imposing temporal uncertainty or requiring self-control for longer periods of time (36) than we used in our task might lead to an increase in ongoing self-control costs.

Self-control research across disciplines suggests that we should be able to induce changes in these costs with changes in motivational state. Consistent with the notion that decisions to engage in self-control arise from a dynamic cost–benefit evaluation (31), we found that willingness to pay for precommitment increased as the cost of failing to adhere to one's diet increased. When faced with losing \$15 in addition to failing to adhere to their diet, participants were willing to pay more to restrict access to temptation; thus, precommitment became more valuable (Study 2 and 4). These findings are consistent with work showing that motivational incentives can increase willingness to engage in self-control strategies (37, 38) and demonstrate overall cost sensitivity in self-control mechanisms.

We also observed an increase in the cost of self-control under stress but only when participants showed differences in neuroendocrine stress responses as measured by a salivary cortisol (Study 3). Stress exposure has long been thought to compromise self-control (15, 16), and this intuition has been borne out in a large body of empirical work—from the cognitive neuroscience literature that shows that stress diminishes cognitive capacity and flexibility (39, 40) and selectively reduces goal-directed control of decisions (41, 42) to the clinical literature in which stressors remain a primary risk factor for the emergence and relapse of addiction-related disorders (43). Our findings provide a direct

test of whether stress compromises self-control by increasing its associated cognitive cost. This relationship was also observed beyond participants assigned stress condition, as higher self-control costs were positively associated with perceived stress scores. We note that in Study 4, where both stress and incentives were imposed before the self-control task, we did not observe an added increase in bids above and beyond that of the Study 2 incentive group. However, this may be due to the fact that we did not see evidence of a cortisol elicitation from the stress manipulation. Future work may seek to use alternate stress induction techniques (e.g., social or cognitive stressors) to test whether other types of stressors lead to additive (or interactive) effects with motivational incentives on self-control costs. Finally, one can imagine decision environments in which stress might impair precommitment decisions given their reliance on prospective thinking. Our task utilized both prospective and experiential exposure to temptation, but it would be interesting to examine how stress affects precommitment decisions that rely fully on future prospection or memory retrieval.

Finally, in Study 5, confirming that our bids did in fact reflect the cost of self-control and not random baseline bidding behavior, we found that the average magnitude of bids scaled with increasing levels of temptation. By testing bidding behavior across varying degrees of temptation level and a broader range of times with the food, we were able to confirm that willingness to pay for precommitment tracks with the increased cost of resisting temptation. One open question is why these costs were higher in Study 5 relative to Studies 1 through 4. One possibility is that these differences stem from the higher hunger ratings reported in Study 5 participants (mean = 9.75, SD = 0.45) relative to those from Studies 1 through 4 (mean = 6.41, SD = 1.78). Another possibility is that these differences arose from the task structure used in Study 5. Since participants bid on a wider range of food items, temptation levels, and time amounts with these foods, participants' perceived self-control costs may be greater for high-tempting foods that are evaluated in the presence of low- and medium-tempting foods. Nonetheless, our findings suggest that measuring control costs can reveal unique information about the subjective experience of resisting temptation that may not be attainable using existing measures of self-control.

Given its importance as a strategy to help individuals achieve their long-term goals by reducing self-control costs, a number of recent studies have begun to measure preferences for precommitment in the laboratory (44–48). For example, Crockett and colleagues provided an empirical demonstration that precommitment facilitates choices for larger, later rewards as opposed to smaller, sooner ones when explicitly offered as a choice strategy in an intertemporal choice task (e.g., viewing erotic images that varied in reward magnitude and delay). This study was the first to demonstrate that the primary neural circuits underlying precommitment (e.g., lateral frontopolar cortex) (46) are distinct from that of standard self-control (“willpower”) tasks that rely on the effortful inhibition of impulses (e.g., dorsolateral prefrontal cortex, inferior frontal gyrus), consistent with the notion that these self-control strategies rely on different neurocognitive processes. These data, coupled with more recent work showing that higher impulsivity and metacognitive awareness leads to stronger preferences for precommitment (47), are consistent with the view that precommitment decisions engage future planning and prospection and are driven by an awareness of subjective self-control costs.

Our results add to this growing literature by demonstrating that individuals not only show preferences for precommitment, but they are willing to pay to adopt such strategies. Our task extends existing studies of precommitment outside of classic intertemporal and effort-based choice paradigms by quantifying the cost of resisting temptation under different conditions. One relevant question for future work is how we can further dissociate the cost–benefit processes that drive precommitment decisions. For example, future research might address to what extent

individuals pay to remove temptation in order to avoid a self-control failure versus to diminish the disutility of resisting temptation irrespective of predicted failures. One recent study, for example, found that precommitment increases motivation to engage in effortful action that leads to larger rewards. In the context of an effort- and delay-based choice task, precommitment choices were found to be driven by a desire to reduce opportunity costs and secure adequate motivation to endure the longer delay or higher effort required to attain larger rewards rather than to avoid a failure of willpower per se (48). Future work may attempt to further dissociate the motivational processes that underlie these decisions and how changes in temptation intensity or temporal uncertainty may alter them.

A number of limitations should be noted for future work. First, we acknowledge that there are likely many different components to what makes the exertion of self-control cognitively costly. For example, there may be cognitive costs to resisting temptation and also personal and health costs associated with self-control failures. While our task does not currently dissociate among the components that feed into self-control costs more generally, this is an open area for future work. Second, unlike some studies in the human stress literature, we included both men and women in each of our samples. However, we did not measure menstrual phase, oral contraceptive use, or cycle-dependent sex hormones, which have been shown to impact stress responses in women (49, 50). Future work measuring such factors may reveal interesting patterns in control costs that we were unable to detect presently. Finally, despite every effort to eliminate them from our procedure (*SI Appendix*), we cannot definitively rule out the possibility that demand characteristics may have contributed to these effects in some way. Our use of an incentive-compatible task meant that we used no deception, misdirection, or foils regarding the goal of the experiment (i.e., measuring self-control costs); thus, we cannot rule out that participants' choice behavior was not in some way influenced by their clear understanding of the study's objective. We note that this would not necessarily explain other features of findings, such as the relationship between bidding behavior and hunger ratings (*SI Appendix, Fig. S4*), perceived stress (Fig. 7A), or dietary length (Fig. 7B). Nonetheless, future work may seek to explore how more stringent mitigation measures for demand effects influence this form of decision-making.

In summary, we report a task for measuring the subjective cost of self-control. Our findings are consistent with emerging work across disciplines suggesting that self-control and its failures can be seen as fundamentally rational responses to a complex world in which individuals' trade-off the cognitive cost of resisting immediate temptation against the benefits of achieving future goals.

Methods

Participants. A total 138 healthy young adult participants that indicated they were on a diet to maintain or lose weight participated in the study (*SI Appendix*). Participants were recruited using flyers posted on and around the New York University campus as well as electronic advertisements on New York University's Department of Psychology website. Participants were excluded prior to participation for the following reasons: 1) pregnancy; 2) high blood pressure or a heart condition; 3) history or medication for neurological or psychiatric disorders; 4) diabetes, metabolic disorders, food allergies, or history of eating disorders; and 5) use of corticosteroids or beta-blockers. All participants provided written informed consent in accordance with experimental procedures approved by the New York University Committee on Activities Involving Human Subjects. All research and experimental procedures were performed in accordance with approved Institutional Review Board guidelines and regulations.

Subjects were paid \$15 per hour plus a \$10 bidding endowment. Six participants from the stress groups were unable to complete the CPT task and were thus excluded. Two additional participants were removed prior to data analysis, because they revealed that they were on special diets to sustain (and ideally increase rather than decrease) weight, and two others were excluded for being on medication (revealed after the experiment ended). Our final

analysis included a total of 128 healthy participants (84 women) with a mean age of 24.37 (SD = 7.07; range = 18 to 55).

General Procedure (Study 1 through 4). Hungry, healthy dieters were asked to abstain from eating 3 to 4 h before participating in the study. Upon arrival at the laboratory, participants provided informed consent and were escorted to the experiment room for a 10-min acclimation period, after which they rated their current hunger level (from 1 to 10; *SI Appendix, Figs. S3 and S4 and SI Appendix, Results*), completed the food-rating and -ranking scales (*SI Appendix, Fig. S5 and SI Appendix, Results*), and provided basic information about their current diets (*SI Appendix, Fig. S6*). After baseline cortisol was collected, participants received their \$10 (cash) endowment. They then received instructions regarding the self-control decision task and were explicitly informed about which high- and low-tempting foods they would be making choices about during the study (participants in Studies 2 and 4 were further instructed that they would lose a \$15 bonus payment provided at the end of the study if they consumed the tempting food at any point). All participants then completed either the CPT or control task and were given a 10-min break in the experiment room before an additional cortisol sample was collected. This break was implemented to ensure that cortisol levels induced by the CPT would begin to peak in coordination with the choice task. Participants then completed the self-control decision task (see *Decision-Making Task*), during which they indicated the maximum amount that they would be willing to pay to remove the high-tempting food from the room and replace it with the low-tempting food. After the 30-min bidding phase was complete, a final cortisol sample was collected, and the final phase of the experiment began, during which participants were required to remain in the experiment room with the high- or low-tempting food for the final 30 min of the study (see *Bid Realization Procedure*). Once this final 30-min phase was complete, participants were paid for their time and left the laboratory.

Stress Induction Technique. All participants in the stress group (Studies 3 and 4) completed the CPT, for which participants submerged their right forearms in ice-cold water (0 °C to 4 °C) for 3 min continuously. All participants in the control group (Studies 1 and 2) followed the exact same procedure using room-temperature water (30 °C to 35 °C). The CPT is widely used in laboratory settings to model the effects of mild-to-moderate stress and reliably generates both autonomic nervous system and HPA axis activation, as measured by increased physiological arousal, neuroendocrine responses, and subjective stress ratings (28–30).

Neuroendocrine Assessment. Saliva samples were collected throughout the study to assess cortisol concentrations, which serve as neuroendocrine markers of stress response. Participants were run between 12 and 5 PM to control for diurnal rhythms of stress hormone levels. Saliva samples were collected using a high-quality synthetic polymer-based salivette placed under participants' tongues for 2 min. Participants were initially seated in the experiment room for a 10-min acclimation period, during which they drank 4 oz of water to clear any residual saliva. Samples were collected at baseline (sample 1), 10 min after CPT/control task administration, which was directly before the self-control decision task (sample 2), and directly before the realization phase of the task began (~30 min after the CPT/control task administration; sample 3). Samples were immediately stored in a sterile tube and preserved in a freezer set to –80 °C. Samples were analyzed by the Psychobiological Research Laboratory of the University of Trier, Germany, using a time-resolved immunoassay with fluorometric detection [DELFLIA; cf. (51)]. Duplicate assays were conducted for each sample, and the average of the two values was used in our analyses. Any samples that contained insufficient saliva could not be analyzed and were excluded from analysis. Cortisol data were log-transformed to account for the skewed nature of cortisol distributions (*SI Appendix, Fig. S1 A and B and SI Appendix, Results*).

Food Item Scales and Selection. In order to select a high- and low-tempting food item for each individual, participants completed a series of food-rating scales prior to the study (*SI Appendix, Fig. S5 and SI Appendix, Results*). Participants separately rated 20 food items (*SI Appendix, Fig. S7*) on how tasty, healthy, and tempting these items were from 1 (not at all) to 10 (very much so). Participants then ranked these 20 food items from best (#1) to worst (#20) for their current diet. Low-tempting foods were identified as those that fell in the lowest 20% of taste and temptation ratings and the highest 20% of health ratings, and that was ranked in the upper 10% of foods best for the participants' current diet. Conversely, high-tempting foods were identified as those that fell in the upper 20% of taste and temptation ratings and the lowest 20% of health ratings, and that was ranked in the lowest 10% of foods worst for the participants' current diet.

Decision-Making Task. To directly examine individuals' subjective cost of self-control, we designed an incentive-compatible decision-making task that measured the monetary costs that participants were willing to incur to avoid temptation on a moment-to-moment basis. There were two phases of the task: 1) a bidding phase, during which participants indicated the maximum they would be willing to pay from a \$10 endowment to remove the high-tempting food from the room and replace it with the low-tempting food, and 2) a realization phase, during which participants sat in the experiment room with the food for the final 30 min of the experiment. Participants in Studies 2 and 4 were further instructed that they would lose a \$15 bonus payment at the end of the study if they consumed the tempting food at any point.

On each trial, participants viewed a computer screen that prompted them to enter the maximum amount that they were currently willing to pay to remove the high-tempting food from the room and replace it with the low-tempting food for the next 30 min. Participants registered their bids with an open-response window by using the mouse to control a sliding bar that ranged from \$0 to \$10 (in \$0.01 increments) and clicking the mouse on their selected bid. Bid trials were presented approximately every 3 min (Studies 2 and 4) or 2 min (in an effort to acquire more precise temporal measurements of bids) for a total of 10 bids for Studies 2 and 4 and 15 bids for Studies 1 and 3. Choices were presented using MATLAB PsychToolBox.

After each bid was received, there was a 2% chance that this bid would be immediately employed in a procedure that would lead to the final 30-min phase in which the high-tempting or low-tempting food would be in the room with the subject. This incentivized participants to bid their true value for removing the high-tempting food since the bidding phase of the task could end on any trial, and only the current bid would be used to determine whether the food was removed for the 30-min final phase. The 2% hazard rate also ensured that the majority of bidding trials would not be realized, allowing us to track the dynamics of how self-control costs change over time with greater exposure to temptation. Finally, this feature of the task allowed us to eliminate any effect of temporal discounting on the sequential bids. To realize bids, we used a standard Becker–DeGroot–Marschak (BDM) auction procedure widely used to reveal maximum willingness to pay (see *Bid Realization Procedure*).

For the initial bidding trial, no food was present in the room. If this initial trial was not realized, the food was brought in the room and remained there until a bid was implemented or the bidding phase concluded. If the trial was realized, the BDM procedure determined whether the food was brought into the room for 30 min. Immediately after each bidding trial, participants were notified as to whether that particular bid would be implemented. If no bid was realized by the end of the 30-min bidding phase, the task transitioned automatically to the final 30-min realization phase.

Self-Control Failures. Any quantity of food consumed during the task was considered a self-control failure. In all cases of eating, participants consumed the entire item with the exception of one participant who consumed half of the snack food (potato chips) presented. Any participants that consumed the food did so during the final 30-min realization phase. To reduce observer effects, participants were alone in the experiment room for the duration of the study, so we were not present to measure the precise time at which they might have eaten the tempting food. Whether or not they consumed the food was revealed after participants left the laboratory.

Bid Realization Procedure. To determine whether the participant won or lost the opportunity to replace the high-tempting food with a low-tempting food, a standard economic auction procedure was implemented (BDM). Participants selected a "selling chip" from a bag (chips ranged from \$0 to \$10 in \$0.01 increments), and this selected chip represented the winning sell price. This randomly selected sell price was then compared against the participant's current bid. If the bid price was greater or equal to that of the sell price, then the participant won the auction, and they paid the sell price (not the bid price) from their endowment to have the high-tempting food removed. They kept any remaining endowment money. If the bid was lower than the sell price, the participant lost the auction. In this case, the high-tempting food remained in the room for the remainder of the experiment, and the participant would keep the entire endowment. This procedure incentivizes participants to report their true maximum willingness to pay.

Study 5. An additional 20 participants were recruited following the same recruitment, screening, informed consent, and payment procedures as Studies 1 through 4 (*Participants*). Two participants were excluded for using medications on our exclusion criteria list. Due to a computer software error, data from two other participants was not recorded. Participants were asked to refrain from eating prior to coming into the laboratory and began by rating the same series of 20 food items (see *SI Appendix, Fig. S7* for choice

set) on how tasty, healthy, and tempting these items were from 1 (not at all) to 10 (very much so). These ratings allowed us to select a low-, medium-, and high-tempting food for each participant. On each trial, participants viewed an image of a snack food that varied on temptation level (low, moderate, or high), quantity (small, medium, or large), and amount of time for which participants would potentially have to spend with the food (1, 3, 5, 10, 15, 20, 25, 30, 45, or 60 min). Participants viewed the food for 4 s, and they entered how much they were willing to pay (from \$0 to \$10) from a \$10 endowment to avoid that food, given the temptation level, quantity, and time amount. After the 90 trials were completed, one trial was randomly selected, and the same BDM auction procedure was used to identify whether the participants won or lost the auction based on their bid for that given trial (*Bid Realization Procedure*). All participants remained in the experimental room for 1 h after the bidding task was complete in order to control for the cost of time (i.e., to ensure bids did not reflect an aversion to waiting for the allocated amount of time relative to the cost of control). If participants lost the auction, the food was present in the room for the amount of time stated on the selected trial (e.g., if the trial depicted 15 min with the food, then the participant spent 15 min of the full hour with the food). If they won the auction, then the food was not present during this amount of time.

1. R. H. Strotz, Myopia and inconsistency in dynamic utility maximization. *Rev. Econ. Stud.* **23**, 165–180 (1956).
2. R. H. Thaler, H. M. Shefrin, An economic theory of self-control. *J. Polit. Econ.* **89**, 392–406 (1981).
3. F. Gul, W. Pesendorfer, Temptation and self-control. *Econometrica* **69**, 1403–1435 (2001).
4. F. Gul, W. Pesendorfer, Self-control and the theory of consumption. *Econometrica* **72**, 119–158 (2004).
5. W. Mischel, E. B. Ebbesen, Attention in delay of gratification. *J. Pers. Soc. Psychol.* **16**, 329–337 (1970).
6. W. Mischel, E. B. Ebbesen, A. R. Zeiss, Cognitive and attentional mechanisms in delay of gratification. *J. Pers. Soc. Psychol.* **21**, 204–218 (1972).
7. W. Mischel, Y. Shoda, M. I. Rodriguez, Delay of gratification in children. *Science* **244**, 933–938 (1989).
8. J. T. McGuire, J. W. Kable, Rational temporal predictions can underlie apparent failures to delay gratification. *Psychol. Rev.* **120**, 395–410 (2013).
9. R. F. Baumeister, E. Bratslavsky, M. Muraven, D. M. Tice, Ego depletion: Is the active self a limited resource? *J. Pers. Soc. Psychol.* **74**, 1252–1265 (1998).
10. M. Muraven, D. M. Tice, R. F. Baumeister, Self-control as limited resource: Regulatory depletion patterns. *J. Pers. Soc. Psychol.* **74**, 774–789 (1998).
11. R. F. Baumeister, T. F. Heatherton, Self-regulation failure: An overview. *Psychol. Inq.* **7**, 1–15 (1996).
12. M. Muraven, R. F. Baumeister, Self-regulation and depletion of limited resources: Does self-control resemble a muscle? *Psychol. Bull.* **126**, 247–259 (2000).
13. R. Kurzban, A. Duckworth, J. W. Kable, J. Myers, An opportunity cost model of subjective effort and task performance. *Behav. Brain Sci.* **36**, 661–679 (2013).
14. A. F. T. Arnsten, Stress signalling pathways that impair prefrontal cortex structure and function. *Nat. Rev. Neurosci.* **10**, 410–422 (2009).
15. G. Hockey, *Stress and Fatigue in Human Performance* (Wiley, 1983).
16. D. H. Holding, "Fatigue" in *Stress and Fatigue in Human Performance*, G. Hockey, Ed. (Wiley, 1983), pp. 145–167.
17. D. Ariely, K. Wertenbroch, Procrastination, deadlines, and performance: Self-control by precommitment. *Psychol. Sci.* **13**, 219–224 (2002).
18. G. Bryan, D. Karlan, S. Nelson, Commitment devices. *Annu. Rev. Econ.* **2**, 671–698 (2010).
19. G. Ainslie, Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychol. Bull.* **82**, 463–496 (1975).
20. T. C. Schelling, Self-command in practice, in policy, and in a theory of rational choice. *Am. Econ. Rev.* **74**, 1–11 (1984).
21. P. C. Fishburn, A. Rubinstein, Time preference. *Int. Econ. Rev.* **23**, 677–694 (1982).
22. A. Tversky, R. H. Thaler, Preference reversals. *J. Econ. Perspect.* **4**, 201–221 (1990).
23. D. Fudenberg, D. K. Levine, A dual-self model of impulse control. *Am. Econ. Rev.* **96**, 1449–1476 (2006).
24. D. Laibson, Golden eggs and hyperbolic discounting. *Q. J. Econ.* **112**, 443–477 (1997).
25. S. Frederick, G. Loewenstein, T. O'Donoghue, Time discounting and time preference: A critical review. *J. Econ. Lit.* **40**, 351–401 (2002).
26. J. W. Kable, P. W. Glimcher, The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* **10**, 1625–1633 (2007).
27. G. M. Becker, M. H. DeGroot, J. Marschak, Measuring utility by a single-response sequential method. *Behav. Sci.* **9**, 226–232 (1964).
28. W. Lovallo, The cold pressor test and autonomic function: A review and integration. *Psychophysiology* **12**, 268–282 (1975).

Data Analysis. All statistical analyses for behavioral and cortisol data were carried out using Statistical Package for the Social Sciences (version 20.0, 2011; IBM). Due to the skewed nature of cortisol concentration distributions documented in the literature (52), cortisol values were log-transformed in order to better approximate a Gaussian (normal) distribution. Data were tested for equal variances using Mauchly's sphericity tests, and Greenhouse–Geisser corrections were performed to address any violations of sphericity. ANOVA with repeated measures was used to analyze all choice (i.e., bidding) and cortisol data. Post hoc comparisons were conducted using Student's *t* tests when appropriate. All tests were two-tailed and considered statistically significant when $P < 0.05$.

Data Availability. Anonymized behavioral choice data have been deposited in Open Science Framework (<https://osf.io/p5942/>).

ACKNOWLEDGMENTS. We thank all members of P.W.G.'s laboratory, especially Anna Konova, Alexandra Mellis, and Christopher K. Steverson, for helpful feedback, comments, and discussion. This work was supported by NIH Grant F32MH110135 and a National Alliance for Research on Schizophrenia & Depression Young Investigator Grant from the Brain & Behavior Research Foundation to C.M.R. and NIH Grant R01DA038063 to P.W.G.

29. A. L. McRae *et al.*, Stress reactivity: Biological and subjective responses to the cold pressor and Trier Social stressors. *Hum. Psychopharmacol.* **21**, 377–385 (2006).
30. M. Velasco, J. Gómez, M. Blanco, I. Rodríguez, The cold pressor test: Pharmacological and therapeutic aspects. *Am. J. Ther.* **4**, 34–38 (1997).
31. E. T. Berkman, C. A. Hutcherson, J. L. Livingston, L. E. Kahn, M. Inzlicht, Self-control as a value-based choice. *Curr. Dir. Psychol. Sci.* **26**, 422–428 (2017).
32. W. Kool, M. Botvinick, Mental labour. *Nat. Hum. Behav.* **2**, 899–908 (2018).
33. A. Westbrook, T. S. Braver, Cognitive effort: A neuroeconomic approach. *Cogn. Affect. Behav. Neurosci.* **15**, 395–415 (2015).
34. A. Shenhar, M. M. Botvinick, J. D. Cohen, The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron* **79**, 217–240 (2013).
35. J. T. McGuire, J. W. Kable, Medial prefrontal cortical activity reflects dynamic re-evaluation during voluntary persistence. *Nat. Neurosci.* **18**, 760–766 (2015).
36. B. Blain, G. Hollard, M. Pessiglione, Neural fatigue and intertemporal choice. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 6967–6972 (2016).
37. M. S. Hagger *et al.*, A multi-lab preregistered replication of the ego-depletion effect. *Perspect. Psychol. Sci.* **11**, 546–573 (2016).
38. R. M. Krebs, C. N. Boehler, M. G. Woldorff, The influence of reward associations on conflict processing in the Stroop task. *Cognition* **117**, 341–347 (2010).
39. D. Schoofs, O. T. Wolf, T. Smeets, Cold pressor stress impairs performance on working memory tasks requiring executive functions in healthy young men. *Behav. Neurosci.* **123**, 1066–1075 (2009).
40. A. Petzold, F. Plessow, T. Goschke, C. Kirschbaum, Stress reduces use of negative feedback in a feedback-based learning task. *Behav. Neurosci.* **124**, 248–255 (2010).
41. A. R. Otto, C. M. Raio, A. Chiang, E. A. Phelps, N. D. Daw, Working-memory capacity protects model-based learning from stress. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 20941–20946 (2013).
42. L. Schwabe, O. T. Wolf, Stress-induced modulation of instrumental behavior: From goal-directed to habitual control of action. *Behav. Brain Res.* **219**, 321–328 (2011).
43. R. Sinha, How does stress increase risk of drug abuse and relapse? *Psychopharmacology (Berl.)* **158**, 343–359 (2001).
44. M. J. Crockett *et al.*, Restricting temptations: Neural mechanisms of precommitment. *Neuron* **79**, 391–401 (2013).
45. J. Schwartz *et al.*, Healthier by precommitment. *Psychol. Sci.* **25**, 538–546 (2014).
46. A. Soutschek *et al.*, Binding oneself to the mast: Stimulating frontopolar cortex enhances precommitment. *Soc. Cogn. Affect. Neurosci.* **12**, 635–642 (2017).
47. A. Soutschek, P. N. Tobler, Know your weaknesses: Sophisticated impulsiveness motivates voluntary self-restrictions. *J. Exp. Psychol. Learn. Mem.* **46**, 1611–1623 (2020).
48. B. Studer, C. Koch, S. Knecht, T. Kalenscher, Conquering the inner couch potato: Precommitment is an effective strategy to enhance motivation for effortful actions. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **374**, 20180131 (2019).
49. C. Kirschbaum, B. M. Kudielka, J. Gaab, N. C. Schommer, D. H. Hellhammer, Impact of gender, menstrual cycle phase, and oral contraceptives on the activity of the hypothalamus-pituitary-adrenal axis. *Psychosom. Med.* **61**, 154–162 (1999).
50. D. H. Hellhammer, S. Wüst, B. M. Kudielka, Salivary cortisol as a biomarker in stress research. *Psychoneuroendocrinology* **34**, 163–171 (2009).
51. R. A. Dressendorfer *et al.*, Synthesis of a cortisol-biotin conjugate and evaluation as a tracer in an immunoassay for salivary cortisol measurement. *J. Ster. Biochem. and Mol. Biol.* **43**, 683–692 (1992).
52. R. Miller *et al.*, Comparison of salivary cortisol as measured by different immunoassays and tandem mass spectrometry. *Psychoneuroendocrinology* **38**, 50–57 (2013).