NIH Public Access
Author Manuscript
J Exp Psychol Gen. Author manuscript; available in PMC 2015 February 01.

Published in final edited form as:
J Exp Psychol Gen. 2014 February ; 143(1): 131–141. doi:10.1037/a0031048.

A labor/leisure tradeoff in cognitive control

Wouter Kool and **Matthew Botvinick**
Department of Psychology and Princeton Neuroscience Institute, Princeton University

Abstract

Daily life frequently offers a choice between activities that are profitable but mentally demanding (cognitive labor) and activities that are undemanding but also unproductive (cognitive leisure). Although such decisions are often implicit, they help determine academic performance, career trajectories, and even health outcomes. Previous research has shed light both on the executive control functions that ultimately define cognitive labor and a 'default mode' of brain function that accompanies cognitive leisure. However, little is known about how labor/leisure decisions are actually made. Here, we identify a central principle guiding such decisions. Results from three economic-choice experiments indicate that the motivation underlying cognitive labor/leisure decision-making is to strike an optimal *balance* between income and leisure, as given by a joint utility function. The results reported establish a new connection between microeconomics and research on executive function. They also suggest a new interpretation of so-called ego-depletion effects, and a potential new approach to such phenomena as mind-wandering and self-control failure.

Imagine a high-school student sitting at her home computer on a school night. Moment by moment, she faces a recurring decision: Should I focus on my calculus homework, or should I take a break to daydream or watch online videos? The student's decision is essentially one between cognitive labor and cognitive leisure, mental work and mental rest. In the language of cognitive psychology, the labor in question involves the effortful engagement of executive control: a set of functions, dependent on the prefrontal cortex, that configure information-processing resources for the execution of computationally intensive, non-routine tasks (Miller & Cohen, 2001). The student's moment-by-moment decision is thus between a gainful activity that demands robust cognitive control and alternatives that do not.

Cognitive labor/leisure decisions play an obvious role in determining academic and career performance. They may also have safety implications in settings such as air-traffic control or power-plant operation; neuroimaging research shows that lapses of attention and response errors are preceded by a shift in activation from dorsal cortical areas underlying cognitive control to a default-mode network that is characteristically most active at rest (Eichele et al., 2008; Weissman, Roberts, Visscher, & Woldorff, 2006). Furthermore, cognitive labor/leisure decisions may go awry in clinical disorders like depression or addiction, leading to maladaptive decision-making (Baler & Volkow, 2006; Murphy et al., 2001; van der Plas, Crone, van den Wildenberg, Tranel, & Bechara, 2008; Wenzlaff & Bates, 1998). Given these and other considerations, it is important to understand how such decisions are made.

Recent behavioral and neuroscientific research indicates that the exertion of cognitive control is intrinsically costly or aversive, and that control is only robustly recruited in the presence of relevant incentives (Jimura, Locke, & Braver, 2010; Kool, McGuire, Rosen, &

All correspondence should be addressed to Wouter Kool, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08540, wkool@princeton.edu.

The authors declare that they have no conflict of interest with respect to their authorship or the publication of this article.

Botvinick, 2010; McGuire & Botvinick, 2010). Cognitive labor/leisure decisions thus appear to involve a cost-benefit analysis, which weighs the payoff from cognitive work against its inherent cost.

Because cost-benefit analyses typically result in binary go/no-go decisions, it is tempting to think of cognitive labor/leisure decisions as involving a categorical, all-or-none choice between labor and leisure. The aim of the present work was to evaluate a more complex but also more interesting account, according to which decision-making is based on preferences for a particular balance or mixture of cognitive labor and leisure.

A formal framework for understanding labor/leisure tradeoffs is provided by economic models of labor supply (Nicholson & Snyder, 2008). Such models address the question of how workers choose their preferred hours: Given a particular hourly wage and a 'budget' of available hours per week, how does a worker choose to allocate time between work and leisure? Classical labor supply theory proposes that workers value both income and leisure, and that overall value or utility is a joint function of both. A graphical depiction of a representative utility function is presented in Figure 1A. Here, the x-axis indicates leisure time, the y-axis total income from labor, and the z-axis utility. The utility function takes the form of a surface, which attaches a subjective value to each pairing of income with leisure time (the lower portion of the figure shows this as a contour plot). A central proposition of labor supply theory is that this utility function is quasi-concave, favoring combinations of labor and leisure that are relatively balanced.

Economic labor supply theory provides a simple and compelling account of how labor/leisure decisions are made, given this form of utility function. Note that the worker's available hours and wage define a range of attainable income-leisure combinations, referred to in labor supply theory as the budget constraint. An example budget constraint is shown in Figure 1A, appearing as a diagonal line in the lower part of the figure and as a grey plane in the surface plot. According to labor supply theory, the worker selects, from among all feasible income/leisure combinations, the one yielding the greatest utility. In Figure 1A, this optimum is marked in red. As the figure illustrates, the concave utility function leads to decisions that favor a balance between income and leisure, resisting a strong bias toward either one or the other

Labor supply theory predicts interesting, non-intuitive shifts in choice behavior with variations in wage. For example, Figure 1B diagrams the effect of an *income-compensated wage decrease*. Here, the worker is given an initial baseline wage (in blue) and selects the income/leisure combination associated with maximal utility (corresponding, in the diagram, to the point of contact between the budget constraint line and the highest isoutility curve with which it intersects). Next, a wage reduction is imposed (in red), but the worker is also given a 'freebie' payment up front. This unearned income precisely compensates for the wage reduction, in the sense that if the worker chooses the same work hours as she did before the wage change, her total income (including the freebie payment) will also remain unchanged (the blue and red budget constraints cross at the initial allocation). Despite the availability of this *status quo ante* income/leisure combination, the geometry of the utility surface in Figure 1A dictates that the worker will reduce her hours, settling for a smaller total income. To see this, note that the new optimal point lies on an indifference curve associated with higher utility than the previous equilibrium.

An opposite but otherwise analogous effect is induced by an *income-compensated wage increase* (Figure 1C). Here, participants are initially given a freebie payment and baseline wage (in blue). In a second session, an increase in wage is paired with a reduction in unearned income (in green). Although the worker once again has access to the same income/

leisure combination that she chose before the wage change, labor supply theory predicts that the worker will sacrifice leisure time in order to attain a larger total income.

A question that often arises when first encountering a description of these effects is, Why bother with the income compensation? Why not consider changes in work allocation under simple increases or reductions in wage? It may appear obvious that labor should increase and leisure decrease in response to an increase in wage. However, as it turns out, this sort of simple monotonic relationship between wage and work is not at all predicted by labor supply theory. It is in fact an important corollary of the theory that increases in wage can, under certain circumstances, yield a reduction in work. Empirical findings support this aspect of labor supply theory. Non-monotonic relations between wage and labor supply have been observed in behavioral-economic studies involving both human (Bigelow & Liebson, 1972) and animal workers (Collier & Jennings, 1969; Green, Kagel, & Battalio, 1987). Income compensation is thus a critical measure, since it yields unambiguous predictions: Labor supply theory strongly predicts that income-compensated wage increases (decreases) should be followed by increases (decreases) in effort supplied.

When such effects are empirically observed, they constitute direct evidence for preferences that balance income against leisure. Income-compensated wage manipulations therefore offer a powerful tool for probing labor/leisure decision-making. In recognition of this, such wage manipulations have been applied in numerous experimental and field studies on labor markets (Charness & Kuhn, 2010; Dickinson, 1999; Fehr & Goette, 2007), as well as in animal conditioning research (Chen, Lakshminarayanan, & Santos, 2006; Kagel, Battalio, & Green, 1995). In the present work, we leveraged the same approach to test for balance-oriented preferences in human cognitive control. Our general hypothesis was that cognitive labor/leisure decisions operate like the labor/leisure decisions studied in labor economics, tending toward a balance between income and leisure. Based on this, our specific prediction was that cognitive labor/leisure decisions should respond to income-compensated wage manipulations as dictated by labor supply theory.

In what follows, we report three experiments. In all three, participants freely allocated time between a highly demanding cognitive task, which paid a wage ('labor'), and a non-demanding task, which did not ('leisure'). The 'leisure' task was closely matched to the 'labor' task on stimulus and response characteristics, demands for physical effort, and intrinsic interest; the primary factor that differentiated the two tasks was cognitive demand, the key ingredient distinguishing cognitive labor from cognitive leisure as we intend these terms. In each experiment, we measured changes in time allocation induced by particular wage manipulations, testing predictions derived from the labor supply theory model (Figure 1B and 1C).

Experiment 1

Our first experiment imposed an income compensated wage reduction. As illustrated in Figure 1B, the labor supply theory account predicts that this manipulation should induce participants to forego a previously selected labor/leisure combination, choosing instead to sacrifice income by reducing time in cognitive labor.

Methods

Participants—Thirty-three participants (23 female) from the Princeton University community provided consent and received course credit or a nominal payment. All procedures were approved by Princeton University Institutional Review Board.

Procedures—Participants were asked to spend forty minutes performing a combination of two computer-based tasks: a *high-demand* task that placed heavy demands on executive function and a very *low-demand* task that did not. Both tasks comprised of the sequential presentation of photographs of humans faces (one face every 1750 msec). In the low-demand task, if this showed a child, the participant was to depress a joystick button (20% of all trials). To minimize demands for cognitive control, a tone would always accompany the presentation of child images. The high-demand task involved identical stimuli and timing (including the tone accompanying child faces), but called for a button-press when any face matched the one presented three steps earlier (also 20% of all trials). Note that this task involves the active maintenance and updating of human faces in working memory, and therefore requires cognitive control (Braver et al., 1997).

Stimuli for the two tasks were presented on opposite sides of a computer screen, and participants chose between the tasks by directing a joystick to the relevant side. Instructions indicated that participants could allocate their time between the tasks however they pleased, and could switch from one to the other whenever and however often they wished. Importantly, participants were paid a wage, denominated in pieces of M&M's candy (Mars Incorporated, Hackettstown, NJ), for each trial dedicated to the high-demand task. The low-demand task, in contrast, yielded no such pay. A number indicating total accrued reward for the present session was always visible at the top of the computer monitor, and increased with each trial presented during the high-demand task according to the size of the wage. Participants were explicitly instructed that payment of the wage was not contingent on the accuracy of their responses on the high-demand task, but only on the number of trials dedicated to this task. To further discourage a special focus on error monitoring, no feedback on response accuracy was provided. Participants were simply instructed, at the outset of the experiment, to do their best on both tasks. Each session of the experiment began with a short practice session, allowing familiarization with both tasks and with the payoff regime, and concluded with delivery of earned income. Because the tone accompanying child faces was also played during the high-demand task, it was possible that participants used this tone to more easily identify target trials with child faces. However, paired *t*-tests on the data of this and the following experiment showed that accuracy was not higher for target trials with child faces than for target trials with adult faces, $ps > 0.10$.

The experiment involved two sessions, separated by one to two weeks. In the first, all participants were accorded the same baseline wage (0.029 pieces of candy for each high-demand trial), and cumulative earnings (*income*) and the time allocated to the high- and low-demand tasks — time in *labor* and (relative) *leisure* — were recorded. In session two, an income-compensated wage reduction was imposed: Participants were explicitly told that the wage from session one was reduced by 50%, but that they also received a 'freebie' candy payment up front (though were asked not to consume any of it until after completion of the experiment). This unearned income was calibrated, separately for each participant, so as to make it possible for the participant to allocate exactly the same amount of time to labor as in session one and accumulate precisely the same total income (see Figure 1B). Our prediction, drawn from labor supply theory, was that the wage change would induce participants to forgo this option, settling instead for a smaller total income along with increased leisure time.

Because the wage level in session one was established *a priori*, without information about each participant's baseline preferences, we anticipated that the wage might induce some participants to allocate the large majority of this session either to labor or to leisure. It is important to note that such events would not contradict any prediction from labor supply theory. Despite the balance principle it implies, labor supply theory does allow that particular incentive regimes may yield an optimum strongly favoring either labor or leisure.

Furthermore, some participants may have considered their time budget (as defined in the introduction) to extend beyond the forty minutes of the testing period, leading to choices that balanced within-session labor against pre- or post-session leisure. Notwithstanding these points, we were concerned that cases of nearly all-or-none time allocation might reduce the sensitivity of our design by introducing a ceiling or floor effect. For this reason, only participants who allocated more than 5% of each session, i.e., two minutes, to both labor and leisure (n = 20) were included in the final data set[1]. Only participants satisfying this criterion in session one were invited to complete session two.

Analyses

The central prediction was that leisure time would increase (and total income therefore decrease) from session one to session two. Time in leisure was compared between sessions one and two by way of a two-tailed paired *t*-test, excluding participants (n = 2) for whom the difference between sessions fell greater two standard deviations beyond the sample mean. Additional analyses, aimed at verifying task compliance, are described in conjunction with results[2].

Results and Discussion

To confirm compliance with task instructions on both high- and low-demand tasks, we computed response accuracies and *d*' indices for each. Results, summarized in Table 2, indicated that participants made an adequate effort to perform both tasks as instructed, responding well above chance levels (single-sample *t*-test, *p*s < 0.001, two tailed). In both sessions, *d*' scores and accuracy were lower for the high-demand task than for the low-demand task (paired *t*-tests, *p*s < 0.001, two tailed), validating our demand manipulation.

Labor/leisure allocations were positively correlated between the two sessions ($r = 0.72$, $p < 0.001$) indicating that preferences were relatively stable over the intersession interval. However, the wage manipulation had a significant effect. Mean leisure time and income for sessions one and two are shown in Figure 2A. On average across participants, the income-compensated wage decrease triggered a 29% increase in leisure time and a 6% decrease in income, a statistically significant shift in the predicted direction ($t(17) = 2.48$, $p < 0.05$, Cohen's $d = 0.58$).

Experiment 2

The results of Experiment 1 supported the idea that effort-based decision making is governed by a preference for balance between control and leisure. As predicted by labor supply theory, an income-compensated wage reduction led participants to reduce time in labor, settling for less income, even though it was possible precisely to recreate an earlier-chosen labor/leisure combination. Experiment 2 tested the complementary prediction, by imposing an income-compensated wage increase. As articulated in Figure 1C, labor supply theory here predicts that participants will here forego a previously selected labor/leisure combination, choosing to increase labor in order to accrue a larger income.

[1]There were eleven participants who allocated more than 95% of the session to the high-demand task and two participants who allocated greater than 95% to the low-demand task.

[2]All *t*-test analyses reported in the present paper were repeated using a non-parametric procedure (signed-ranks test), yielding similarly significant results in every case.

Methods

Participants—Twenty-one participants (15 female) members of the Princeton University community participated, providing informed consent and receiving course credit or a nominal payment.

Design and procedures—The procedure was largely identical to Experiment 1. However, here, in session one, participants were accorded a lower baseline wage and received an up-front freebie payment. In session two, the wage was doubled and participants received an up-front payment smaller than the payment in session one. This new freebie payment was calibrated so as to make it possible for each participant to allocate exactly the same amount of time to labor as in session one and to accumulate precisely the same total income (see Figure 1C). Again, our prediction from labor supply theory was that the wage change would induce participants to forgo this option, causing them instead to sacrifice leisure time in favor of a larger total income.

As in Experiment 1, measures were taken to mitigate floor and ceiling effects. Here, this was accomplished by focusing analysis on participants (n = 16) who allocated at least 5% of each session to both labor and leisure tasks[3]. In contrast to Experiment 1, all participants completed both sessions one and two.

Analyses

The central prediction was that mean leisure time would decrease (and total income thus increase) from session one to session two. This was tested by a two-tailed paired *t*-test, excluding participants (n = 1) for whom the difference between sessions fell greater two standard deviations beyond the sample mean. Additional analyses are once again described in conjunction with results.

Results and Discussion

As summarized in Table 2, response accuracies again indicated compliance with task instructions, far exceeding chance levels (single-sample *t*-test, $p < 0.001$, two-tailed). As in Experiment 1, participants showed better performance for the low-demand task than for the high-demand task (paired *t*-tests, *p*s < 0.001, two tailed).

Again as in Experiment 1, time-allocation was positively correlated across sessions one and two ($r = 0.81$, $p < 0.001$). But once again, the wage manipulation induced significant changes in task choice. Mean leisure time and income for sessions one and two are shown in Figure 2B. Across participants, the income-compensated wage increase triggered a 13% decrease in leisure time and a 9% increase in income, once again a statistically significant shift in the predicted direction ($t(14) = 2.36$, $p < 0.05$, Cohen's $d = 0.61$). This effect remained significant even when the analysis included participants who dedicated less than 5% of one session to either labor or leisure ($t(19) = 3.29$, $p < 0.004$, Cohen's $d = 1.51$).

Experiments 1 and 2: Additional analyses

One potential concern in interpreting the time-allocation results from Experiments 1 and 2 is that participants might have based their choices exclusively on the proffered wage, ignoring unearned income. As noted in the introduction, it is not necessarily the case that simple (uncompensated) increases in wage yield increases in work. Indeed, in the context of cognitive tasks, past empirical work has failed to reveal any simple relationship between

[3]There were four participants who allocated more than 95% of the session to the high-demand task and one participant who allocated greater than 95% of the time to the low-demand task.

incentive magnitude and effort investment, as reflected in task performance (Bonner & Sprinkle, 2002; Camerer & Hogarth, 1999). Nevertheless, it is worth asking whether the shifts in effort allocation observed in Experiments 1 and 2 might reflect a reaction to simple wage changes alone, in isolation from unearned income. To evaluate this, we compared leisure time between Experiment 1, session one and Experiment 2, session two, sessions that involved equal wages but different freebie payments (pair *a, d* in Figure 2). We also compared leisure time in Experiment 1, session two against leisure time in Experiment 2, session one, since here again wages were equal but unearned income differed (pair *b, c* in Figure 2). In both cases, leisure time differed across the relevant sessions (two sample *t*-tests, *p*s < 0.05). This indicates that time allocation was not based on wage alone. Rather, participants appear to have based their labor/leisure decisions jointly on wage and unearned income.

We interpret the results of Experiments 1 and 2 in terms of effort costs, or the intrinsic cost of cognitive control. It is partly relief from these costs, we propose, that underlies the inherent 'utility of leisure,' as it is termed in labor supply theory. However, one might speculate that what is aversive in the high-demand task is not its cognitive demand, *per se*, but rather its association with frequent response errors. Of course, our participants were given no explicit feedback on response accuracy, and knew that wage delivery was not contingent on performance. More importantly, data from another study (Kool et al., 2010) show that human decision makers avoid cognitive demand even in the absence of error commission. Error avoidance thus seems unlikely to be the only factor underlying the present results. However, to evaluate the potential role of errors, we conducted two additional analyses. First, based on the recommendation of an anonymous reviewer, we tested whether the wage effects from Experiments 1 and 2 would remain significant after controlling for shifts in response accuracy on the high-demand task from session 1 to session 2. Repeated-measures ANCOVAs, using inter-session change in d' as a covariate, revealed intact effects of wage on leisure time for both experiments. Second, within each experimental session, we examined whether participants tended to switch away from the high-demand task after the occurrence of errors. We considered trials where the participant switched from the high- to the low-demand task (*switch* trials) and trials where the participant remained on the high-demand task (*stay* trials), comparing between these the mean accuracy over the preceding run of high-demand trials. Repeated-measures ANOVAs revealed no significant difference in either session of Experiments 1 and 2 for history lengths up to eleven trials. At longer history lengths (up to thirty trials), with no correction for multiple comparisons, an effect appeared in only one of four sessions (Experiment 2, session 2, history lengths 12, 13 and 16–30). In sum, taken together with previous research, the data do not compellingly support an account based entirely on error avoidance. Further evidence along these same lines is presented in conjunction with the next experiment.

Experiment 3

Experiments 1 and 2 provide support for an income/leisure tradeoff in cognitive control allocation, analogous to the one described in economic labor supply theory. However, one might suspect that the tradeoff evident in those experiments arises as an *emergent* feature of cognitive control dynamics. An influential line of research has suggested that cognitive control is subject to fatigue, or resource depletion (Muraven, Tice, & Baumeister, 1998). If this is correct, then perhaps participants in Experiments 1 and 2 allocated time between tasks by simply responding, moment by moment, on the basis of their current level of fatigue, resting when highly 'depleted,' and returning to the high-demand task when 'replenished.'

If this were the whole story, then the effects observed in our first two experiments, during on-line effort allocation, should disappear when allocation decisions must be made prospectively. The present experiment examined effort allocation in precisely this setting.

Methods

Participants—Nineteen participants (12 female) from the Princeton University community participated, providing informed consent and receiving either course credit or a nominal payment.

Design and Procedures—The experiment started with two minutes of practice on both the high- and low-demand tasks used in Experiments 1 and 2. Next, prior to entering a twenty-minute task-performance period, participants used a graphical interface to allocate time prospectively between the high- and low-demand tasks (Figure 3). Each participant made 32 prospective allocation decisions, involving a variety of wages and up-front payments, on the understanding that one of their choices would be chosen randomly and used to program stimulus presentation during the later task-performance period. Embedded in the choice task was a set of randomly interleaved trial-pairs. One member of each pair served as a baseline, while the other implemented an income-compensated wage decrease or increase. On half of all trials, the awarded candy pieces were M&M's, and on the other half Reese's Pieces (The Hershey Company, Hershey, PA). Upon completion, participants received the amount of candy selected on a randomly chosen time-allocation trial. Wage effects were quantified as the average change in work-time between a baseline trial and the associated wage-change trials. We predicted that both wage manipulations would again cause participants to forego their baseline time-allocation choices, leading them to sacrifice income in the case of wage decreases, and to sacrifice leisure in the case of wage increases.

Analyses are described in conjunction with results. Tests on wage effects excluded subjects (n = 1) falling more than two standard deviations from the sample mean, on either wage increase or decrease trials. All *t*-tests were two-tailed.

Results and Discussion

Table 2 lists mean response accuracies and d' scores for the low- and high-demand tasks, which as in previous studies were well above chance ($ps < 0.001$), indicating task compliance.

Mean leisure time and income before and after wage changes are shown in Figure 2C (for wage decreases) and 2D (increases). Across participants and trials, income-compensated wage decreases triggered a 29% increase in leisure and a 12% decrease in income ($t(17) = 3.41$, $p < 0.05$; Cohen's $d = 0.80$), while income-compensated wage increases triggered a 15% decrease in leisure and a 11% increase in income ($t(17) = 2.79$, $p < 0.05$, Cohen's $d = 0.66$).

Interestingly, the 'elasticity' of time allocation – that is, its sensitivity to a unit change in wage – appeared to differ reliably across participants. The degree to which each participant reduced work in the face of wage reductions correlated with the degree to which they increased work in response to a boost in wages ($r = 0.68$, $p < 0.001$; Figure 4). Further analyses suggested that floor or ceiling effects were not wholly responsible for this correlation; the size of wage-change effects did not significantly correlate with mean leisure-time allocation ($r = 0.31$, $p = 0.20$), or with the absolute deviation of leisure-time allocation from 50% ($r = -0.39$, $p = 0.11$).

Beyond its primary goals, the present study also provided an opportunity to address a lingering concern from Experiments 1 and 2. Under *Additional Analyses*, we compared cases where wage was matched but effort allocation differed, taking the difference as evidence that participants took unearned income into account. However, the sessions compared in that analysis differed not only in terms of unearned income. They also differed in terms of context: One of the sessions (the one drawn from Experiment 2) was preceded by a session involving a lower wage, while the other (session 1 from Experiment 1) was not. The difference in effort allocation could thus, arguably, have arisen from an effect of context on the evaluation of wage. Specifically, the wage offered in Experiment 2, session 2, might have appeared more rich and motivating, given the contrast with the earlier, lower wage. In order to evaluate this possibility, as well as to further validate the role of unearned income, we conducted a linear regression for each participant in Experiment 3, modeling the chosen leisure time on each trial as a function of (1) the present wage, (2) the present unearned income, and (3) the difference between present wage and the wage offered on the immediately preceding trial (Table 3). We tested whether each of these regressors, on average, explained variance in participants' choice behavior. Leisure time was, not surprisingly, predicted by present wage ($p < 0.01$). More importantly, it was also predicted by unearned income ($p < 0.05$), but not reliably predicted by wage difference ($p = 0.41$). This pattern of results further supports the integral role of income compensation in driving the main effects from all three studies.

In order to rule out a key role for error avoidance, we leveraged the data from the initial practice session. We reasoned that if effort allocation were based entirely on error rates rather than cognitive demand, then leisure time on baseline trials should correlate across participants with accuracy on the high-demand task during the practice block. However, no such correlation was observed, regardless of whether accuracy was quantified as d', proportion correct, or proportion correct only on target trials ($p > 0.80$ in all cases).

General Discussion

To summarize, we found across three experiments that income compensated wage-changes impacted cognitive labor/leisure decisions in a fashion consistent with economic labor supply theory: Even when participants had the option to recreate a previously selected combination of income and cognitive leisure, wage reductions led them to work less and settle for a smaller income, and wage increases led them to give up leisure in favor of a larger income. The observed pattern of behavior suggests that cognitive labor/leisure decisions are guided by preferences that jointly weigh income and leisure.

Our findings directly parallel results from both experimental labor economics (Charness & Kuhn, 2010; Fehr, Goette, & Zehnder, 2009) and animal conditioning research (Chen et al., 2006; Conover & Shizgal, 2005; Kagel et al., 1995), indicating that the principles that guide effort allocation both in labor markets and in animal foraging may also apply in the more rarified domain of human cognitive control.

On the broader stage, effects of incentives on cognitive effort and performance have been of interest both in basic-science settings (Brehm, Wright, Solomon, Silka, & Greenberg, 1983; Kool et al., 2010; Ryan & Deci, 2000b; Sarter, Gehring, & Kozak, 2006), and in applied settings including education (Ryan & Deci, 2000a; Weiner, 1985). Our work suggests that, within such work, it may be fruitful to apply formal tools from labor economics. For example, economic research has demonstrated non-monotonic effects of incentives on worker effort (resulting in what is known as the 'backward-bending' labor supply curve), and interactions between wage incentives and social preferences on worker behavior (Camerer, Babcock, Loewenstein, & Thaler, 1997). The present findings open up the question of whether effects such as these might also characterize decision-making in cognitive control.

In addition to work on incentive effects, the present findings also relate closely to research on the cost of cognitive control. A number of recent studies (Botvinick, 2007; Botvinick, Huffstetler, & McGuire, 2009; Kool et al., 2010; McGuire & Botvinick, 2010) have lent empirical support to the long-standing proposition that the exertion of executive control is intrinsically costly or aversive. This principle is evident, for example, in choice tasks where participants are free to select between two lines of action associated with different levels of cognitive demand. Human decision-makers in this setting favor the low-demand option, a demand-avoidance effect that (like some of the present findings) is not completely attributable to ancillary factors such as error avoidance or time on task (Kool et al., 2010).

The demand-avoidance effect fits into the framework we have been developing here via a simple change of sign: In cognitive labor/leisure decisions, the utility of leisure derives, in important part, from the relief it offers from costly control. Equivalently, one may view the cost of control as an opportunity cost, relating to foregone leisure. The beauty of the labor supply model is that it renders the choice between these two perspectives moot. The shape of the utility function says it all.

Additional studies of the demand-avoidance effect have shown that it is reduced in the presence of countervailing incentives, indicating that the inherent cost of control is weighed against potential rewards (Kool et al., 2010). The present work provides a richer picture of this cost-benefit analysis. By default, one might assume a simple linear model, under which the net value of a course of action is computed by subtracting the cost of control from associated rewards. The present findings suggest, instead, that the marginal cost of control varies as a function of context: A unit increment in effort carries a greater subjective cost when one is already working hard than when one is hardly working.

This way of characterizing our results reveals a natural relationship with research on the phenomenon of ‘ego depletion.’ The term refers to the observation that voluntary cognitive effort tends to decline after bouts of forced cognitive exertion (Baumeister, 2002; Baumeister, Bratslavsky, Muraven, & Tice, 1998). According to one influential theory, this effect reflects a depletion of blood glucose, a metabolic resource that is hypothesized to be necessary for the exertion of cognitive control (Baumeister, Vohs, & Tice, 2007; Gailliot et al., 2007). The present work suggests an alternative view of the ‘depletion’ effect, according to which it arises from a particular form of cost-benefit analysis: In the context of prolonged, obligatory mental effort, the marginal cost of further effort is elevated, leading in some cases to a subsequent withdrawal from cognitively challenging activity. By characterizing the depletion effect in this way, the present work situates this specific phenomenon within a larger framework for understanding labor/leisure decisions.

Unlike the glucose depletion account, the present theory is primarily descriptive rather than mechanistic or normative. Our goal has been to account for behavior in terms of an underlying utility function, the same strategy that is used by theories of economic choice (e.g., prospect theory) in approaching phenomena such as loss aversion or risk seeking (Kahneman & Tversky, 1979). This descriptive orientation may appear to place the present work at a disadvantage, when compared with the more mechanistic resource-depletion model. In response to this potential challenge, two points may be offered. First, it should be noted that empirical evidence for the role of glucose (or any other metabolic resource) in the behavioral ‘depletion’ effect is, at best, mixed (Kurzban, 2010), undermining depletion’s force as a mechanistic argument. Second, a range of findings – for example, the disappearance of ‘depletion’ effects in the face of increased incentives (Muraven &

Slessareva, 2003) – has led even proponents of the ego-depletion theory to acknowledge the need for an additional layer, at which effort expenditure is shaped by motivational factors (Hagger, Wood, Stiff, & Chatzisarantis, 2010). The present account, which is pitched precisely at this level, thus fills a gap left even by supposedly mechanistic approaches.

Another appealing aspect of the ego-depletion theory is its normative flavor. The theory appears to make clear why it may be *rational* to withdraw cognitive control in certain contexts: Such withdrawal allows preservation or replenishment a limited metabolic resource, viz., glucose (Muraven, Shmueli, & Burkley, 2006). The present account, as a descriptive theory, does not center on (or depend upon) a normative argument. Having said this, it is possible, if one is willing to speculate, to fit the present descriptive theory within a normative framework. From a computational point of view, there are in fact two reasons why the occasional withdrawal of control may be rational, and why it might thus be rational for decision-making to place a cost on control, as we have proposed. First, a wealth of research indicates that controlled information processing is highly capacity-limited. Thus, a motivational bias in favor of automatic processing would have the effect of reserving this limited computational bandwidth for operations that are likely to have a high payoff (see Kool et al., 2010). Second, and more speculatively, it seems plausible that the 'leisure' state arising when control is disengaged may allow the decision-maker to survey or explore available activities, preventing a myopic focus on one potentially suboptimal line of behavior (Kool et al., 2010). In this sense, the labor/leisure tradeoff demonstrated in our experiments may address a high-level instantiation of the exploitation-exploration tradeoff that arises in reinforcement-learning models of reward-based decision-making (Cohen, McClure, & Yu, 2007) and of foraging (Kolling, Behrens, Mars, & Rushworth, 2012).

We began this article with an imagined 'real-life' example of labor/leisure decision making. The question may now be asked, How do the results of our experiments relate to everyday life? Two aspects of our experimental approach call for comment, in this connection. First, the decisions made by participants in our experiments were deliberate and explicit, a feature that allowed us to interpret them as reflecting participants' actual preferences. Second, the labor/leisure decision in our experiments involved selecting between two active tasks, a measure that allowed us to match 'labor' and 'leisure' conditions in terms of stimulus and response characteristics, physical effort demands, and intrinsic task interest. Some labor/leisure decisions in everyday life clearly involve a deliberate choice between two active tasks. Our earlier homework-versus-video example fits this description, as do certain cases involving 'failures of self-control,' where one throws off the burden of self-discipline (read: cost of control) in order to indulge in a formerly resisted temptation (Kool, McGuire, Wang, & Botvinick, submitted). At the same time, intuition suggests that many everyday labor/leisure decisions are unreflective or even unconscious, and pertain to the level of engagement in a single task rather than to a choice between tasks. The phenomenon of 'mind-wandering,' in particular, seems relevant here, since it is conventionally understood as involving a non-deliberate disengagement from a single, focal task. The theoretical framework we have proposed does plausibly extend to this scenario. On a labor-supply account, mind-wandering would be understood as an implicit, unreflective choice to engage in cognitive leisure (here, avoidance of the focal task in favor of non-demanding but otherwise unspecified mental activities), triggered by an elevation of marginal control costs above task payoffs. Mind-wandering has been shown to reflect a disengagement of cognitive control (McVay & Kane, 2012), and, interestingly, to be accompanied by increased activity in the default-mode brain network (Christoff, Gordon, Smallwood, Smith, & Schooler, 2009). However, further experimentation will be required to establish whether the labor-supply framework we have introduced indeed applies to this and other variants of cognitive labor/leisure decision-making.

What neural mechanisms might underlie the pattern of labor/leisure decision making observed in our experiments? The exertion of cognitive control is well known to engage a specific set of areas within dorsolateral prefrontal cortex, dorsomedial frontal cortex, the insula and parietal cortex (Duncan, 2010). Recent evidence indicates that portions of this network also index, in tandem with related subcortical structures, both the potential payoffs (Kouneiher, Charron, & Koechlin, 2009; Locke & Braver, 2008) and the subjective cost of control (Botvinick et al., 2009; McGuire & Botvinick, 2010). As noted earlier, a parallel literature points to a neural correlate of cognitive 'leisure,' linking the disengagement of cognitive control with activation of a default mode network, centering on ventral prefrontal cortex (Kelly, Uddin, Biswal, Castellanos, & Milham, 2008; Raichle et al., 2001). We speculate that the labor/leisure tradeoff demonstrated in the present behavioral studies reflects the operation of a neural mechanism regulating the balance of activity between these dorsal 'task-positive' and ventromedial 'task-negative' systems (Menon & Uddin, 2010; Sridharan, Levitin, & Menon, 2008). If this is correct, then the formal framework we have introduced here might usefully inform future neuroscientific research into such a mediating neural mechanism.

Acknowledgments

This research was supported by National Institute of Mental Health Grant MH062196 and a grant from the John Templeton Foundation. We also thank N. I. Córdova, J. Kaplan and J. Prabhakar for research assistance.

References

Baler RD, Volkow ND. Drug addiction: the neurobiology of disrupted self-control. Trends in Molecular Medicine. 2006; 12:559–566. [PubMed: 17070107]

Baumeister RF. Ego depletion and self-control failure: An energy model of the self's executive function. Self and Identity. 2002; 1:129–136.

Baumeister RF, Bratslavsky E, Muraven M, Tice DM. Ego depletion: is the active self a limited resource? Journal of Personality and Social Psychology. 1998; 74:1252–1265. [PubMed: 9599441]

Baumeister RF, Vohs KD, Tice DM. The strength model of self-control. Current Directions in Psychological Science. 2007; 16:351–355.

Bigelow G, Liebson I. Cost factors controlling alcoholic drinking. The Psychological Record. 1972; 22:305–314.

Bonner SE, Sprinkle GB. The effects of monetary incentives on effort and task performance: theories, evidence, and a framework for research. Accounting, Organizations and Society. 2002; 27:303–345.

Botvinick MM. Conflict monitoring and decision making Reconciling two perspectives on anterior cingulate function. Cognitive, Affective, & Behavioral Neuroscience. 2007; 7:356–366.

Botvinick MM, Huffstetler S, McGuire JT. Effort discounting in human nucleus accumbens. Cognitive Affective & Behavioral Neuroscience. 2009; 9:16–27.

Braver TS, Cohen JD, Nystrom LE, Jonides J, Smith EE, Noll DC. A parametric study of prefrontal cortex involvement in human working memory. Neuroimage. 1997; 5:49–62. [PubMed: 9038284]

Brehm JW, Wright RA, Solomon S, Silka L, Greenberg J. Perceived difficulty, energization, and the magnitude of goal valence. Journal of Experimental Social Psychology. 1983; 19:21–48.

Camerer CF, Babcock L, Loewenstein G, Thaler R. Labor supply of New York City cabdrives: One day at a time. The Quarterly Journal of Economics. 1997; 112:407–441.

Camerer CF, Hogarth RM. The effects of financial incentives in experiments: A review and capital-labor-production framework. Journal of Risk and uncertainty. 1999; 19:7–42.

Charness G, Kuhn P. What can labor economists learn from the lab? IZA Discussion Paper no. 4941. 2010

Chen MK, Lakshminarayanan V, Santos LR. How basic are behavioral biases? Evidence from capuchin monkey trading behavior. Journal of Political Economy. 2006; 114:517–537.

Christoff K, Gordon AM, Smallwood J, Smith R, Schooler JW. Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. Proceedings of the National Academy of Sciences. 2009; 106:8719.

Cohen JD, McClure SM, Yu AJ. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. Philosophical Transactions of the Royal Society B: Biological Sciences. 2007; 362:933.

Collier G, Jennings W. Work as a determinant of instrumental performance. Journal of Comparative and Physiological Psychology. 1969; 68:659–662.

Conover KL, Shizgal P. Employing labor-supply theory to measure the reward value of electrical brain stimulation. Games and Economic behavior. 2005; 52:283–304.

Dickinson DL. An experimental examination of labor supply and work intensities. Journal of Labor Economics. 1999; 17:638–670.

Duncan J. The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. Trends in Cognitive Sciences. 2010; 14:172–179. [PubMed: 20171926]

Eichele T, Debener S, Calhoun V, Specht K, Engel A, Hugdahl K, von Cramon D, Ullsperger M. Prediction of human errors by maladaptive changes in event-related brain networks. Proceedings of the National Academy of Sciences. 2008; 105:6173–6178.

Fehr E, Goette L. Do workers work more if wages are high? Evidence from a randomized field experiment. American Economic Review. 2007; 97:298–317.

Fehr E, Goette L, Zehnder C. A behavioral account of the labor market: The role of fairness concerns. Annual Review of Economics. 2009; 1:355–384.

Gailliot MT, Baumeister RF, DeWall CN, Maner JK, Plant EA, Tice DM, Brewer LE, Schmeichel BJ. Self-control relies on glucose as a limited energy source: Willpower is more than a metaphor. Journal of Personality and Social Psychology. 2007; 92:325–336. [PubMed: 17279852]

Green L, Kagel JH, Battalio RC. Consumption-leisure tradeoffs in pigeons: Effects of changing marginal wage rates by varying amount of reinforcement. Journal of the Experimental Analysis of behavior. 1987; 55:133–143. [PubMed: 2037823]

Hagger MS, Wood C, Stiff C, Chatzisarantis NLD. Ego depletion and the strength model of self-control: A meta-analysis. Psychological Bulletin. 2010; 136:495–525. [PubMed: 20565167]

Jimura K, Locke HS, Braver TS. Prefrontal cortex mediation of cognitive enhancement in rewarding motivational contexts. Proceedings of the National Academy of Sciences. 2010; 107:8871–8876.

Kagel, JH.; Battalio, RC.; Green, L. Economic choice theory: An experimental analysis of animal behavior. Cambridge: Cambridge University Press; 1995.

Kahneman D, Tversky A. Prospect theory: An analysis of decision under risk. Econometrica: Journal of the Econometric Society. 1979; 47:263–292.

Kelly AM, Uddin LQ, Biswal BB, Castellanos FX, Milham MP. Competition between functional brain networks mediates behavioral variability. Neuroimage. 2008; 39:527–537. [PubMed: 17919929]

Kolling N, Behrens TEJ, Mars RB, Rushworth MFS. Neural mechanisms of foraging. Science. 2012; 336:95–98. [PubMed: 22491854]

Kool W, McGuire JT, Rosen ZB, Botvinick MM. Decision making and the avoidance of cognitive demand. Journal of Experimental Psychology: General. 2010; 139:665–682. [PubMed: 20853993]

Kool W, McGuire JT, Wang GJ, Botvinick MM. A role for effort costs in self-control and intertemporal choice. (submitted).

Kouneiher F, Charron S, Koechlin E. Motivation and cognitive control in the human prefrontal cortex. Nature Neuroscience. 2009; 12:939–945.

Kurzban R. Does the brain consume additional glucose during self-control tasks? Evolutionary Psychology. 2010; 8:244–259. [PubMed: 22947794]

Locke HS, Braver TS. Motivational influences on cognitive control: Behavior, brain activation, and individual differences. Cognitive Affective, & Behavioral Neuroscience. 2008; 8:99–112.

McGuire JT, Botvinick MM. Prefrontal cortex, cognitive control, and the registration of decision costs. Proceedings of the National Academy of Sciences USA. 2010; 107:7922–7926.

McVay JC, Kane MJ. Drifting from slow to “d’oh!”: Working memory capacity and mind wandering predict extreme reaction times and executive control errors. Journal of Experimental Psychology: Learning, Memory, and Cognition. 2012; 38:525–549.

Menon V, Uddin LQ. Saliency, switching, attention and control: a network model of insula function. Brain Struct Funct. 2010; 214:655–667. [PubMed: 20512370]

Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. Neuroscience. 2001; 24:167–202.

Muraven M, Shmueli D, Burkley E. Conserving self-control strength. Journal of Personality and Social Psychology. 2006; 91:524–537. [PubMed: 16938035]

Muraven M, Slessareva E. Mechanisms of self-control failure: Motivation and limited resources. Personality and Social Psychology Bulletin. 2003; 29:894–906. [PubMed: 15018677]

Muraven M, Tice DM, Baumeister RF. Self-control as limited resource: regulatory depletion patterns. Journal of Personality and Social Psychology. 1998; 74:774–789. [PubMed: 9523419]

Murphy FC, Rubinsztein JS, Michael A, Rogers RD, Robbins TW, Paykel ES, Sahakian BJ. Decision-making cognition in mania and depression. Psychological Medicine. 2001; 31:679–693. [PubMed: 11352370]

Nicholson, W.; Snyder, CM. Microeconomic theory: Basic principles and extensions. Mason: Cengage Learning; 2008.

Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, Shulman GL. A default mode of brain function. Proceedings of the National Academy of Sciences. 2001; 98:676–682.

Ryan RM, Deci EL. Intrinsic and extrinsic motivations: Classic definitions and new directions. Contemporary Educational Psychology. 2000a; 25:54–67. [PubMed: 10620381]

Ryan RM, Deci EL. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. American Psychologist. 2000b; 55:68–78. [PubMed: 11392867]

Sarter M, Gehring WJ, Kozak R. More attention must be paid: the neurobiology of attentional effort. Brain Research Reviews. 2006; 51:145–160. [PubMed: 16530842]

Sridharan D, Levitin D, Menon V. A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. Proceedings of the National Academy of Sciences. 2008; 105:12569–12574.

van der Plas EAA, Crone EA, van den Wildenberg WPM, Tranel D, Bechara A. Executive control deficits in substance-dependent individuals: A comparison of alcohol, cocaine, and methamphetamine and of men and women. Journal of Clinical and Experimental Neuropsychology. 2008; 31:706–719. [PubMed: 19037812]

Weiner B. An attributional theory of achievement motivation and emotion. Psychological Review. 1985; 92:548–573. [PubMed: 3903815]

Weissman DH, Roberts KC, Visscher KM, Woldorff MG. The neural bases of momentary lapses in attention. Nature Neuroscience. 2006; 9:971–978.

Wenzlaff RM, Bates DE. Unmasking a cognitive vulnerability to depression: how lapses in mental control reveal depressive thinking. Journal of Personality and Social Psychology. 1998; 75:1559–1571. [PubMed: 9914666]

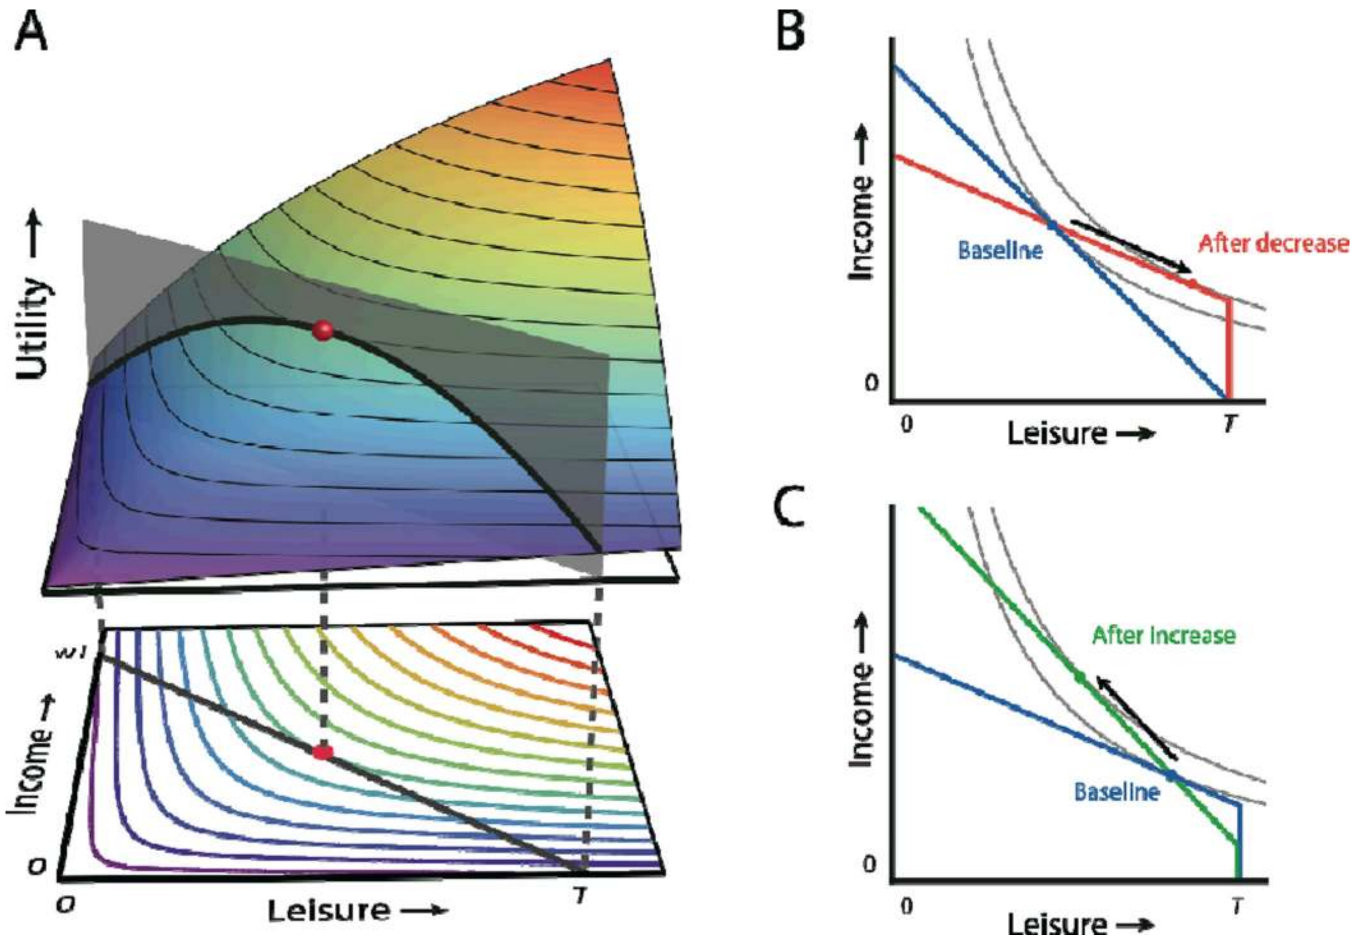

1.

A. Above: A concave utility surface attaching a scalar value to each pairing of income with leisure time. Shown in black are isoutility or indifference curves. Below: A contour map, projected from the indifference curves above. The diagonal line represents a budget constraint, covering the range of income-leisure combinations open to a worker facing *T* available hours and a wage *w*. The grey rectangle above projects this constraint line back up through the three-dimensional utility surface, with the intersection forming an arc. In both top and bottom plots, the red marker indicates the income-leisure combination with the highest utility in the feasible set. B: An income-compensated wage decrease. The initial, baseline wage, together with the total available time *T*, defines the initial budget constraint line (blue). The worker chooses the point along this line that maximizes utility (blue point, with associated indifference curve shown in grey). The wage reduction results in a more horizontal budget constraint (diagonal red line). However, it comes along with unearned income (vertical red line), which causes the new budget constraint to intersect the income/leisure combination originally chosen (blue point). Although this combination is thus still available after the wage change, the worker also has access to a higher-utility combination (red point, with associated isoutility contour) involving increased leisure and reduced income. C. An income-compensated wage increase, illustrated in the same fashion. Here an initial unearned income (vertical blue line) plus an initial wage together define a budget constraint (diagonal blue line). When the wage is increased, the unearned income is reduced (vertical green line), making the originally chosen income-leisure combination (blue point) still available. However, the worker now has access to a higher-utility combination involving less leisure and greater income (green point).

NIH-PA Author Manuscript
NIH-PA Author Manuscript

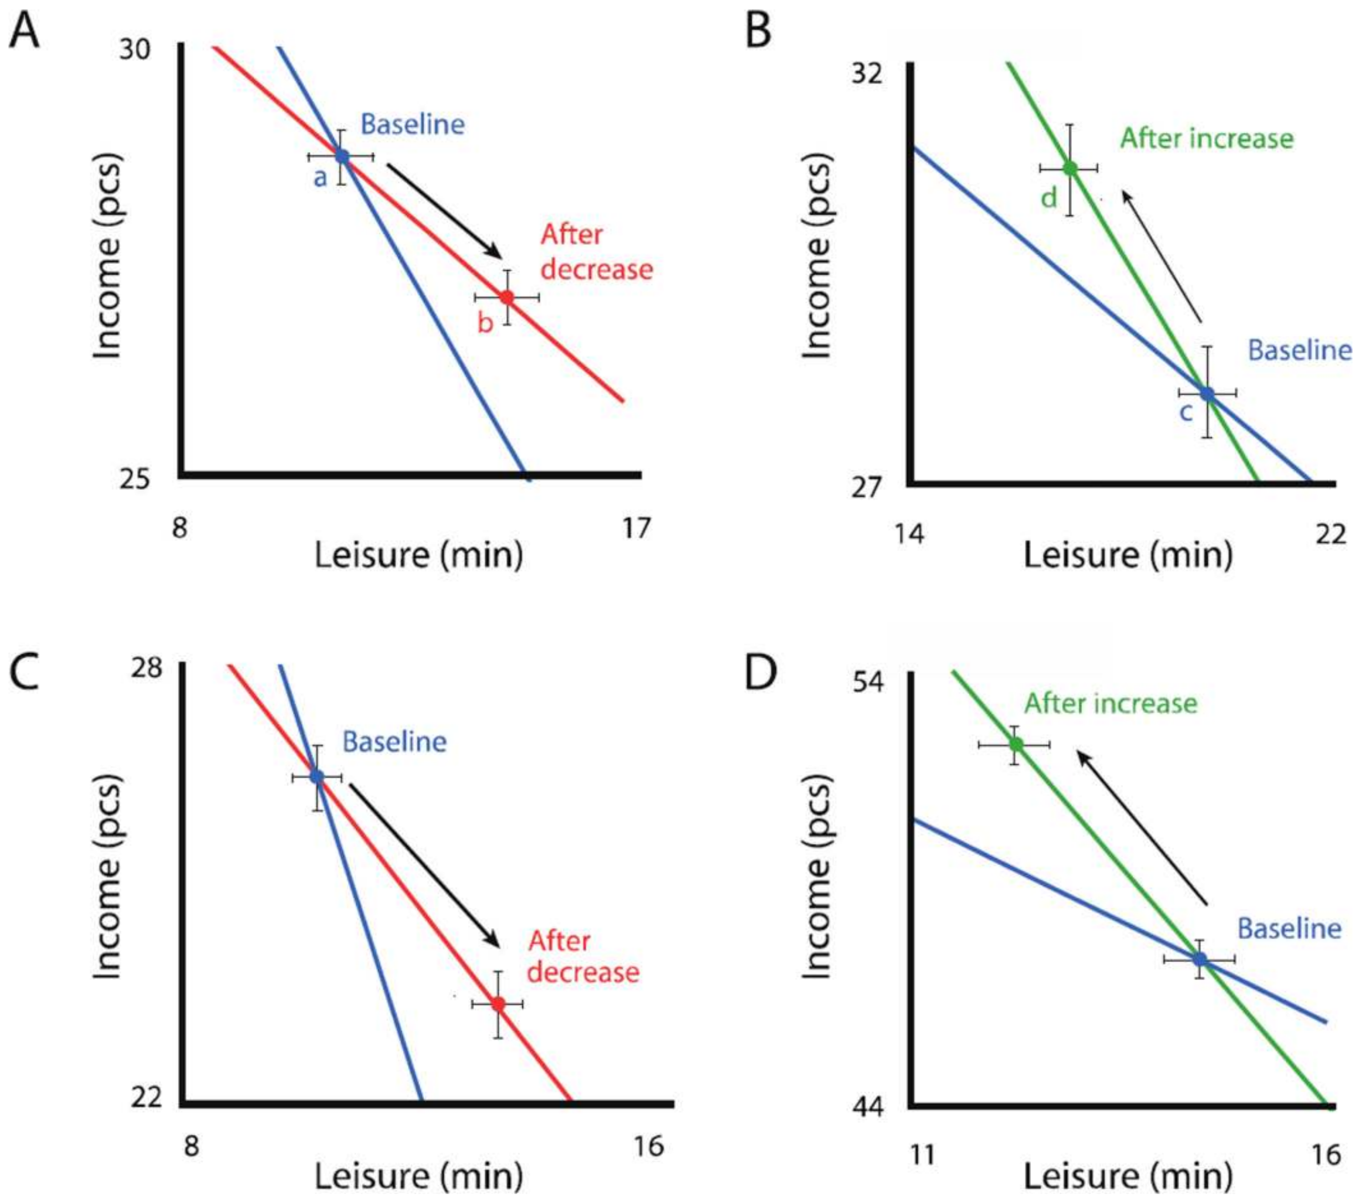

2.

A. Results of Experiment 1 (Compare Figure 1B). Leisure time is in minutes and income in pieces of candy. Red and blue lines indicate mean budget constraints. Error bars indicate within-subject SEM (note that horizontal error bars are simply a scaled version of the vertical bars). B. Results of Experiment 2 (Compare Figure 1C). C. Results for wage-decrease trial-pairs in Experiment 3 (Compare Figure 1B). D. Results for wage-increase trial-pairs in Experiment 3 (Compare Figure 1C). Note that the higher baseline income in this condition reflects relatively high unearned income (see Table 1).

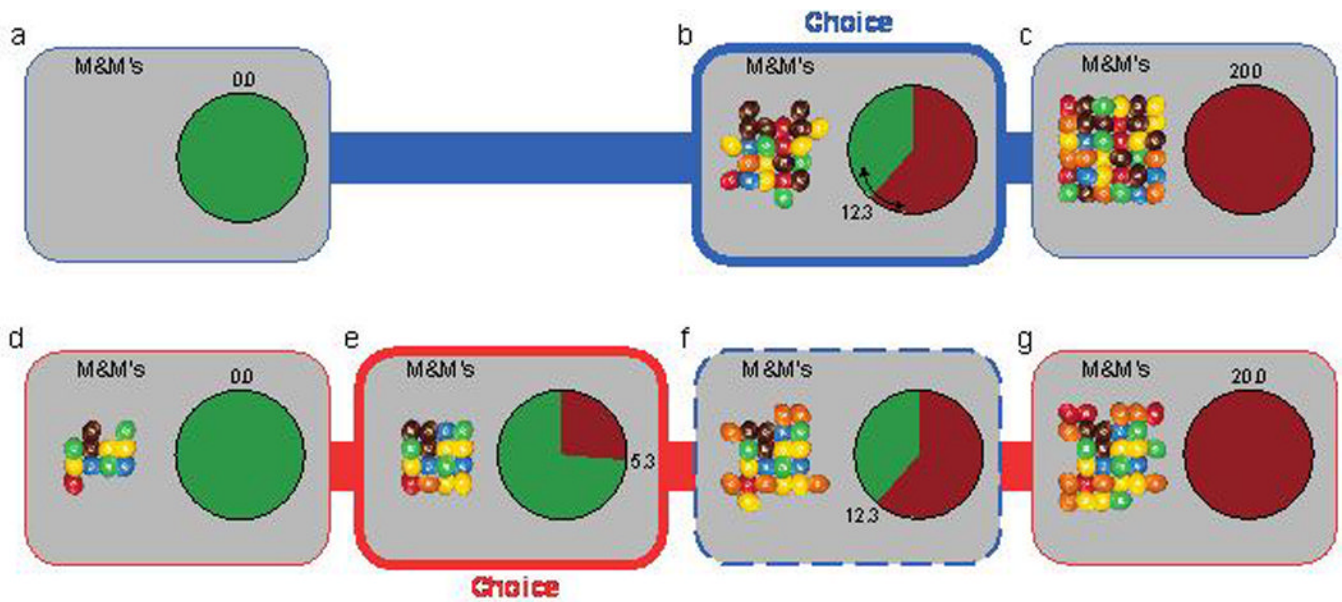

3.

In the choice phase of Experiment 3, the graphical interface showed a single pie chart indicating labor (red) and leisure (green) time, along with an associated total income (in pieces of candy). Participants explored the range of available income/leisure combinations, using arrow keys to increase or decrease leisure time in small increments (curved arrow, not shown in user interface). Panels *a* and *c* in the figure show the extreme allocations options on one choice trial, and *b* the combination actually chosen by the participant. In the lower row, panels *d* and *g* show extreme allocation options on another trial faced by the same participant, which implemented an income-compensated wage decrease relative to the trial diagrammed in the top row. Although the income-leisure combination from panel *b* was available to the participant on this trial (panel *f*), the participant chose a combination involving greater leisure and less income (panel *e*), consistent with the predictions of labor supply theory.

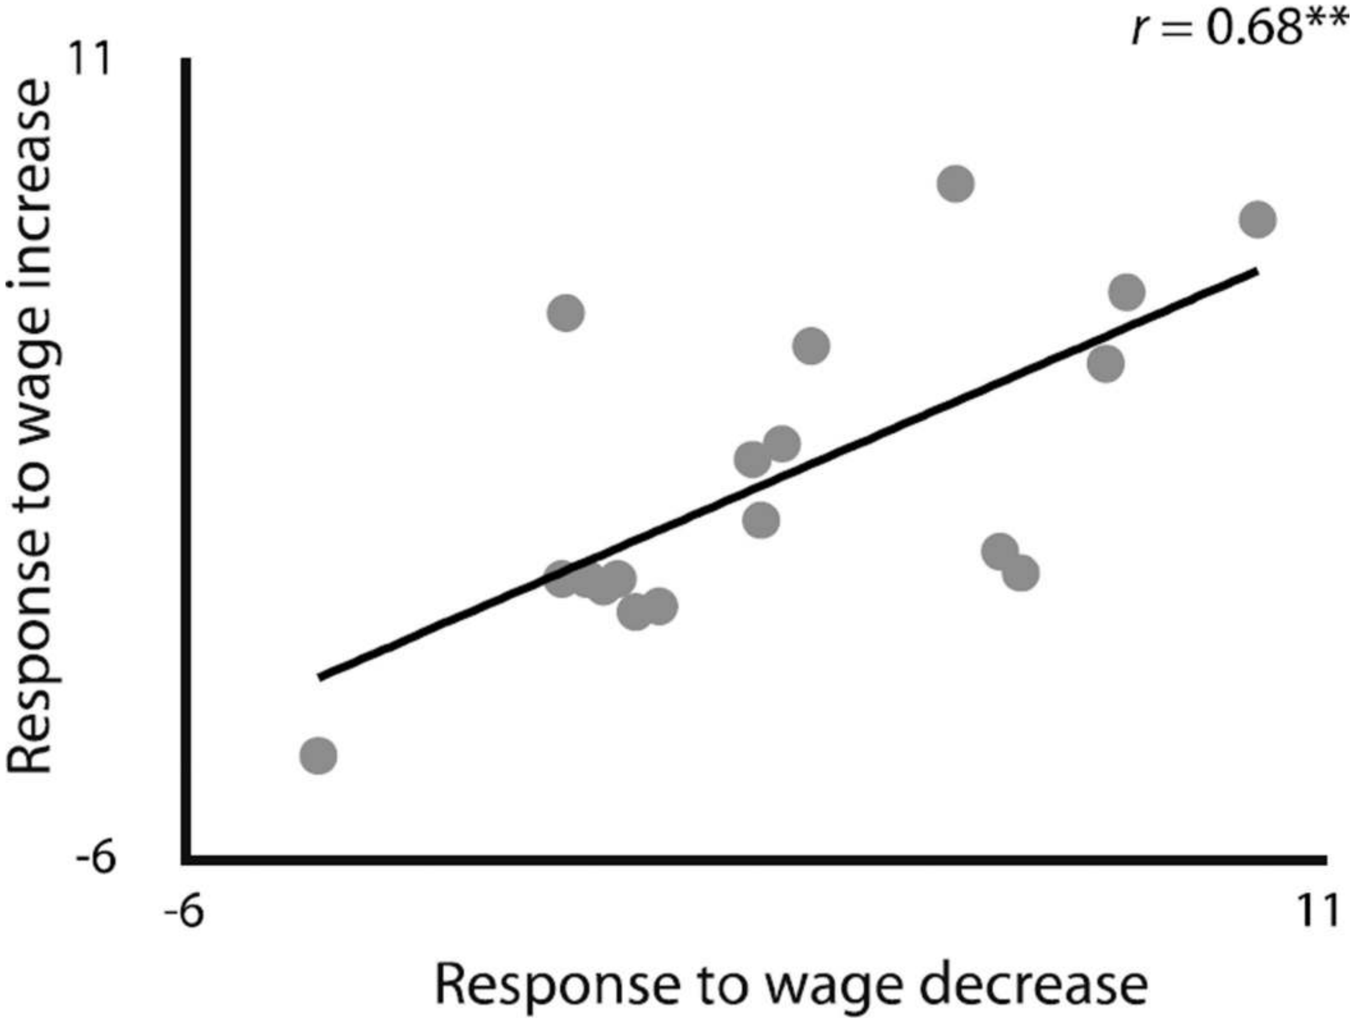

4.
Individual-subject data from Experiment 3. The depicted correlation indicates that, when participants showed a large shift to leisure after a compensated wage decrease, they also showed a large shift to labor after a compensated wage increase. There was significant variability in this measure of elasticity, or the sensitivity to a unit change in wage across participants.

Table 1

Trial-type parameters from Experiment 3. Trials were divided into four groups, each containing four specific trial-types. Each group included one parameterization that served as a baseline condition for the others (first row in each group). The remaining three trial-types each implemented an income-compensated wage decrease (groups A-B) or increase (C-D), with respect to the relevant baseline. Each trial-type formed the basis for four actual trials for each participant. Wages are designated in units of candy-pieces per cumulative minute of labor.

Group	Wage	Unearned Income
A	3.6	None
	3.6 × 1/2	Compensatory (> 0)
	3.6 × 1/3	Compensatory (> 0)
	3.6 × 1/4	Compensatory (> 0)
B	1.8	None
	1.8 × 2/3	Compensatory (> 0)
	1.8 × 1/2	Compensatory (> 0)
	1.8× 2/5	Compensatory (> 0)
C	1.2	48
	1.2 × 1.5	Compensatory (< 48)
	1.2 × 2	Compensatory (< 48)
	1.2 × 3	Compensatory (< 48)
D	0.6	36
	0.6 × 2	Compensatory (< 36)
	0.6 × 3	Compensatory (< 36)
	0.6 × 4	Compensatory (< 36)

NIH-PA Author Manuscript
NIH-PA Author Manuscript

Table 2

Average performance scores (standard error) for the high-demand and low-demand task in the sessions of all three experiments. The table shows the d-prime sensitivity score, average percent correct and reaction times.

	Experiment 1		Experiment 2		Experiment 3
	Session 1	**Session 2**	**Session 1**	**Session 2**	**Practice**
D-prime					
Low-demand	4.17 (0.10)	4.06 (0.10)	4.28 (0.15)	3.60 (0.26)	1.69 (0.11)
High-demand	1.82 (0.08)	1.88 (0.13)	1.63 (0.10)	1.63 (0.11)	3.44 (0.04)
Accuracy					
Low-demand					
Target (%)	98% (0.4)	96% (0.9)	95% (2.2)	85% (5.3)	100% (0)
Non-target (%)	99% (0.2)	99% (0.1)	99% (0.1)	99% (0.3)	99.6% (0.2)
High-demand					
Target (%)	58% (3.2)	54% (4.5)	49% (4.0)	42% (4.7)	61% (4.9)
Non-target (%)	94% (0.8)	96% (0.6)	95% (0.4)	96% (0.4)	92% (1.7)
Reaction time (ms)					
Low-demand	549 (24)	552 (23)	599 (32)	674 (51)	n/a
High-demand	753 (20)	789 (30)	824 (28)	844 (38)	n/a

Table 3

Regression coefficients and statistics for the regression analysis performed on the data of Experiment 3. For each participant, we modeled the time allocated to labor on each trial as the function of the current wage, current unearned income and the difference in wage between the current and the previous trials.

	Regressor Current Unearned Income	Current Wage	$\Delta Wage_{current-previous}$
Regression coefficient (std error)	−0.04 (0.02)	1.34 (0.42)	0.12 (0.15)
t-statistic (df)	−2.22 (17)	3.20 (17)	0.85 (17)
p-value	< 0.05	< 0.01	0.41