Waiting to Choose*

Alex Imas† Michael A. Kuhn‡
Carnegie Mellon University University of Oregon

Vera Mironova§
University of Maryland

November 6, 2017

Abstract

We study the impact of deliberation time on intertemporal choice. Using studies in an online labor market, the laboratory, and a field experiment in the Democratic Republic of Congo, we show that the introduction of waiting periods –designed to prompt deliberation by temporally separating news about choice sets from choices themselves– cause a shift towards substantially more patient decisions. These results cannot be captured by standard models of exponential discounting, nor behavioral models of present bias. Rather, we demonstrate that our findings are consistent with recent work in economics and psychology that identifies the unique role of deliberation in intertemporal choice. We also compare the effect of a waiting period to that of shifting the intertemporal allocation to the future, which separately identifies the impacts of deliberation and present bias. Our results highlight the role of deliberation in decision-making and have implications for policies aimed at reducing myopic choices.

JEL Classification: D03, D90, C91, C93

Keywords: time preferences, waiting periods, heuristics, deliberation, field experiments

*We want to sincerely thank Nageeb Ali, Aislinn Bohren, Michael Callen, Andrew Caplin, Shane Frederick, Xavier Gabaix, Yoram Halevy, David Huffman, Matthew Knepper, David Laibson, George Loewenstein, Paul Smeets, Charles Sprenger, Lise Vesterlund, and seminar participants at Case Western University, Central European University, Harvard University, Maastricht University, Nowegian School of Economics, Ohio State University, Simon Frasier, University of Pittsburgh, University of Arkansas, University of Virginia, the Triennial Choice Symposium, Security and Human Behavior Annual Workship, SITE Experimental Workshop, ECBE Conference, NBER National Security Working Group Winter Meeting for valuable comments and advice.

†Carnegie Mellon University, Social and Decision Sciences, 5000 Forbes Ave, Porter Hall 208, Pittsburgh, PA 15213 USA. E-mail: aimas@andrew.cmu.edu. Web page: www.aleximas.com

‡University of Oregon, Department of Economics, 1285 University of Oregon, Eugene, OR 97403 USA. E-mail: mkuhn@uoregon.edu. Web page: pages.uoregon.edu/mkuhn

§University of Maryland, Department of Government and Politics, 3140 Tydings Hall, College Park, MD 20742 USA. E-mail: vmironov@umd.edu. Web page: vmironova.net

1 Introduction

A large empirical literature on intertemporal choice documents surprisingly nearsighted, myopic behavior. Across a variety of domains, many people turn down large future rewards in favor of smaller, sooner rewards, and postpone unpleasant tasks despite substantial future costs. The classical model of exponential discounting assumes that the discount factor between any two periods is constant. Interpreting these choices within that model leads to a wide range of estimated discount factors. Occasionally, these estimates fall within the bounds of normative plausibility, i.e. close to market rates (Moore and Viscusi, 1990). But more often they are anomalous – either unrealistically low (e.g. Kirby et al. (1999)) or non-constant (e.g. Thaler (1981)).

Observations of anomalously-nearsighted behavior led to the development of theory that relaxes some assumptions of exponential discounting in order to better explain the data. Models of hyperbolic and quasi-hyperbolic discounting (Laibson, 1997; O'Donoghue and Rabin, 1999) capture myopia by allowing for the discount factor to vary depending on the time horizon, with greater discounting between periods closer to the present than further in the future. Other models distinguish between 'fast,' automatic decisions and those based on 'slow,' deliberative processing (Thaler and Shefrin, 1981; Kahneman, 2003). In these models, automatic decisions are characterized by the use of simple heuristics that reduce the complexity of the choice but may result in systematic biases, whereas deliberation leads to decisions that are closer to the predictions of the exponential model. Myopic, nearsighted behavior has been attributed to 'fast' choices and the use of heuristics (Rubinstein, 2003; Read et al., 2013), while more farsighted behavior has been attributed to deliberative processing (Metcalfe and Mischel, 1999).

This literature has typically maintained the assumption that observed choices between sooner-smaller and larger-later rewards are driven by stable underlying preferences – either through a fixed set of discount factors or a fixed propensity to use heuristics (e.g., Rubinstein, 2013; Ericson et al., 2015). In contrast, this paper provides evidence that features of the choice environment can systematically affect an individual's intertemporal choices. Specifically, we show that waiting periods – which are designed to prompt deliberation by temporally separating information about a choice set and the choice itself – can shift individuals' preferences over intertemporal allocations and lead to substantially more patient choices. Through a series of studies, we show that the propensity to

make myopic choices does not seem to be a stable feature of underlying preferences, as individuals make systematically more patient decisions after being prompted to deliberate through a waiting period. This evidence provides support for theory that explicitly considers the role of deliberation time in intertemporal choice (Gabaix and Laibson, 2017).

Waiting periods are often imposed in practice when myopia and impulsivity are perceived to be particularly harmful. Many U.S. states impose a waiting period of up to 14 days between the purchase and receipt of a gun, and recent work has found them to be effective for reducing homicides by as much as 17%, or 750 gun homicides per year (Luca et al., 2017).[1] Waiting periods are also often imposed for those seeking to get married or divorced, and are prescribed as a strategy for avoiding myopic choices in negotiations (Brooks, 2015) and conflict resolution (Burgess, 2004). These policies are predicated on the idea that inserting a delay between when a choice set first comes into focus and when the choice can actually be made may prompt a shift towards more deliberative thinking, and lead to a change in the final decision. Despite the frequent use of waiting periods in these settings, little economic research has been done to examine whether such prompts for deliberation actually affect intertemporal choice.

For example, suppose an individual learns of a substantial windfall in her tax refund. Typically, the waiting time to receive a refund is three weeks after filing, during which she can deliberate what portion to save and what to spend. Now imagine that a firm offered to deliver the refund immediately: would removing the waiting period change the individual's decision to spend or save?[2] The standard model of exponential discounting represents the preference for sooner rewards over later rewards using a constant between-period discount factor, $\delta < 1$. At the moment the refund arrives –irrespective of whether it does so before or after a waiting period– the individual faces the same intertemporal tradeoff. The model predicts that, all else equal, intertemporal choices should be unaffected by waiting periods. Behavioral models that relax the assumption of constant discounting between periods make a similar prediction.[3]

[1] Additionally, Koenig and Schindler (2016) find that the ten states and the District of Columbia that have waiting periods for the purchase of firearms (ranging from 24 hours in Illinois to 14 days in Hawaii) experienced smaller upticks in firearm purchases than other states following the Sandy Hook massacre. Edwards et al. (2016) find that mandatory waiting periods reduce firearm suicides.

[2] This example is pertinent in light of the recent partnership between Walmart, a retail store, and Jackson Hewitt, a tax preparation service. Customers are able to file their taxes at Walmart, with Jackson Hewitt providing the refund immediately on a gift card through a no-cost Refund Anticipation Loan.

[3] This assumes that individuals cannot make binding choices over future allocations and/or naiveté. These predictions

This paper presents evidence that waiting periods have a significant effect on intertemporal choice across three studies. We examine several potential mechanisms, including present bias, and a general tendency to be more prudent when prompted. Our findings highlight an important role of deliberation in intertemporal choices. In the first study, individuals allocated labor and leisure across two hour-long work periods within a single session. We used an online labor market to recruit a population that was experienced in making intertemporal labor-leisure decisions in a similar context. Each participant chose how to allocate real-effort tasks between two work periods, where any time not spent on the effort task could be used to engage in other activities of their choosing. Delaying tasks to a later work period resulted in a greater total task requirement, while choosing to allocate tasks to the earliest available period minimized total work time. In one treatment, participants were given the information about the allocation decision and had the opportunity to make their choice directly after; in another treatment, they were given the same information, but could only make their choice after a one-hour waiting period. Both sets of participants faced the exact same choice set – allocating effort tasks between two work periods – and neither was constrained in the time allowed to make a decision. The only difference is that the latter group faced a waiting period before being able to make a choice, while the former was given the opportunity to make a choice directly after learning the information.

We find that introducing a waiting period had a significant and substantial effect on intertemporal allocations. After the waiting period, participants allocated significantly more effort tasks to the earlier work period, reducing their overall workload as a result. The magnitude of the effect was sizable: 17% of the mean in the treatment without a waiting period, or half of a standard deviation. To rule out the possibility that this effect was driven by differences in the timing of work periods (e.g. one group getting used to waiting, or exogenous information shocks during the delay), we also implemented a control treatment in which the first hour did not involve work, as in the waiting period treatment, but information about the allocation decision was only presented after this delay. If our results were being driven by timing of the work periods, then participants in this control treatment would have made similar allocation decisions to those in the waiting period treatment. In contrast, if waiting periods affect behavior by prompting deliberation over the choice set, allocations in the control treatment should resemble those in the treatment without a waiting

are presented more formally in Section 2.

period. We find strong evidence for the latter.

A potential mechanism driving our results could be that individuals have non-constant discounting (e.g. present bias), are sophisticated about it, and can mentally commit to a plan. We test for this explanation explicitly by implementing a treatment in which participants were given the opportunity to make an allocation decision between two future work periods – the same two as those in the waiting period condition – directly after being informed about the choice set. If behavior in the waiting period treatment could be explained by non-constant discounting, then the allocation decision over two future work periods in this treatment should be similar to the choices of those who faced a waiting period. However, participants in this treatment still allocated significantly more effort tasks to the later work period than those who faced a waiting period before the allocation decision. In our setting, waiting periods have a larger effect on selecting to do more tasks sooner than shifting intertemporal allocations to the future.

Our second study sought to replicate the initial findings, rule out alternative explanations, and further explore the mechanism driving the observed effects. It was run in the laboratory using a similar protocol as the first study. To determine whether social interactions or access to other informational sources were responsible for the impact of waiting periods, the second study enforced tight control over the environment and restricted any participant communication with others. We replicate the effect of waiting periods in this new environment both in sign and magnitude. We also use the laboratory study to test two additional mechanisms that could be driving these results. One potential explanation for our findings comes from the literature on dual-systems processing (Kahneman, 2011). In the basic framework, reliance on one system is characterized by 'fast,' automatic processing while the other system is associated with more deliberative decision-making. In our context, waiting periods may be prompting people to switch from using the 'fast' system to the more deliberative system, leading to more prudent decision-making across domains.

On the other hand, recent work in psychology (Gilbert and Wilson, 2007; Schacter et al., 2007) and economics (Gabaix and Laibson, 2017) consider the role of 'prospection' in intertemporal choice. In these models, people lack perfect information about utility from future outcomes and mentally simulate them to get more accurate representations of the consequences. As a result, deliberation leads to a systematic shift in intertemporal choices, but the effect should be domain-specific; decisions in domains not being considered should remain unaffected. The theoretical

framework of Gabaix and Laibson (2017) makes an additional prediction, that individuals with greater information processing capacities should initially exhibit higher returns from deliberation, and hence respond more dramatically to the introduction of waiting periods in our setting. We find that the waiting period effect is indeed domain-specific, only influencing choices being considered, as well as suggestive evidence that it has a differential impact depending on people's information processing capacities.

We ran a third study in the Democratic Republic of Congo as a proof-of-concept, illustrating how waiting periods can be applied in the field to affect intertemporal decisions and demonstrating the generalizability of our results across contexts. We worked with a small neighborhood grocery store in Bukavu, Democratic Republic of Congo (DRC). Upon arriving at the grocery store, customers received a coupon that could be exchanged for one bag of flour on a pre-specified redemption date. For each day the coupon was saved after this date, its value increased by an additional bag of flour (up to five bags total). In one treatment, customers had the opportunity to redeem the coupon on the same day upon which it was received; in the other, they had to wait one day before being able to redeem the coupon (the value-accrual schedule was thus delayed by a day as well). Thus, the treatment variation enforced a waiting period between coupon receipt and the ability to use it. We find that the introduction of a waiting period led to a significant and meaningful reduction in the fraction of individuals redeeming their coupon on the earliest possible date, for the smallest amount of flour –from 25% without a waiting period to 9% with one.

Our work contributes to the related literature in economics has used the distinction between automatic, heuristic thinking and deliberative processing to explain anomalies in the cross-section of stock returns (Barberis et al., 2013) and excessive focus on the leftmost odometer digits in used-car sales (Lacetera et al., 2012). Rees-Jones and Taubinsky (2016) argue that reliance on heuristics leads to significant misperceptions of the US Federal Income Tax code, while Kessler et al. (2017) show that deliberative processing leads people to choose more efficient allocations in social dilemmas. Dai and Fishbach (2013) show that, similar to our findings, informing participants about a decision one month in the future leads to more patient choices over lotteries involving money or durable goods. However, given the length of the delay, one cannot separately identify the effects of deliberation from the arrival of decision-relevant information during this period (e.g. learning through communication with others), potential opportunities for arbitrage, or changes in

preferences orthogonal to the effects of deliberation (e.g. getting acclimated to process of waiting). Lastly, our findings add to the literature that considers the role of information processing in the interpretation of choice behavior (Caplin et al., 2011; Caplin and Dean, 2015).

The rest of the paper is organized as follows. Section 2 describes both effort allocation experiments, outlines the hypotheses and presents the results. Section 3 presents the application of waiting periods in the field and the results. Section 4 discusses the findings and concludes.

2 Waiting Periods and Effort Allocation

2.1 Online Labor Market

2.1.1 Design and Implementation

Participants (N=122) were recruited from an online labor market (Amazon Mechanical Turk) to complete a series of real-effort tasks over a span of approximately three hours for a $20 payment. The participants were informed that in order to complete the study and earn their payment, they must finish a number of tasks over the course of two one-hour work periods — WP1 and WP2, respectively. We used a single session design to circumvent issues of differential transaction costs and future uncertainty that may affect intertemporal decisions (Benhabib et al., 2010; Cohen et al., 2016). We adopt the approach of Augenblick et al. (2015) in allowing participants to allocate the effort tasks between the two work periods using a series of discretized Convex Time Budgets (Andreoni et al., 2015). Participants could not advance from a work period to the next phase of the study until the full time allocated to that work period elapsed, even if all the allocated effort tasks were completed. They had to remain at or near the computer until whatever time remaining in the period elapsed.[4] The span between finishing the effort tasks and the start of the next work period was explicitly labeled as free time, during which the participants could engage in any activities of their choice.[5] Therefore, the decision to allocate tasks between work periods

[4] After the period elapsed, the participant had to continue either working on the next set of tasks or complete the end line survey; otherwise the experiment would expire and she would not be paid.

[5] We did not want to restrict participants' ability to choose their preferred activity, and recommended that they could spend this time watching streaming movies, reading a book, etc. End-line surveys suggest participants indeed spent the time on leisure activities. Responses included "I cooked and watched television," "I read a book on my Nook. I've had a very stressful day and it was nice to have some free time to do that," "I mostly just listened to music and read some articles. All responses available upon request.

Table 1: Choices in the Online Study

Budget	Max. WP1 Tasks	Max. WP2 Tasks	# of Options	Interest Rate
1	40	60	11	50%
2	40	50	11	25%
3	40	45	6	12.5%
4	40	40	11	0%

WP1 and WP2 refers to Work Period 1 and 2, respectively. Maximum tasks allocated to one work period imply that zero tasks would be allocated to the other work period. The last column lists implied one-hour interest rates.

involved an intertemporal tradeoff between work and free time, which we refer to as leisure from now on.

The task was designed to be onerous and effortful, consisting of counting the numbers of zeros in large, randomly generated table of zeros and ones (Falk et al., 2006; Abeler et al., 2011). Pretests revealed that each 10x15 table took roughly one minute to complete. Participants encountered the tables one at a time and could not advance until they entered the correct answer. Before they were presented with information about the effort allocation budgets and decisions, participants had to successfully complete two sample tasks in order to become familiar with the task.

After completing the sample tasks, participants were informed that they would face a series of choices to allocate effort tasks between WP1 and WP2. One choice was drawn at random and implemented as the actual work requirement. Each participant made allocation decisions using four convex time budgets that varied in the implied interest rate for putting tasks off to the later work period. Every budget allowed for the possibility of doing 40 tasks in WP1. For example, Budget 1 offered the the possibility of 40 tasks in WP1 and no tasks in WP2, no tasks in WP1 and 60 tasks in WP2, or any of nine evenly-spaced convex combinations of those extremes. Implied interest rates varied by budget, from 50% for Budget 1 to 0% for Budget 4. Table 1 presents the convex budgets.

Participants were also given two binary choices that served as manipulation checks for sensitivity to the interest rate. The first offered a choice between 40 tasks in WP1 (and zero tasks in WP2) or 35 tasks in WP2 (and zero tasks in WP1), a negative interest rate; the second offered a choice between 40 tasks in WP1 (and zero tasks in WP2) or 41 tasks in WP2 (and zero tasks in WP1), a very small positive interest rate close to zero. The goal of these questions was to determine the

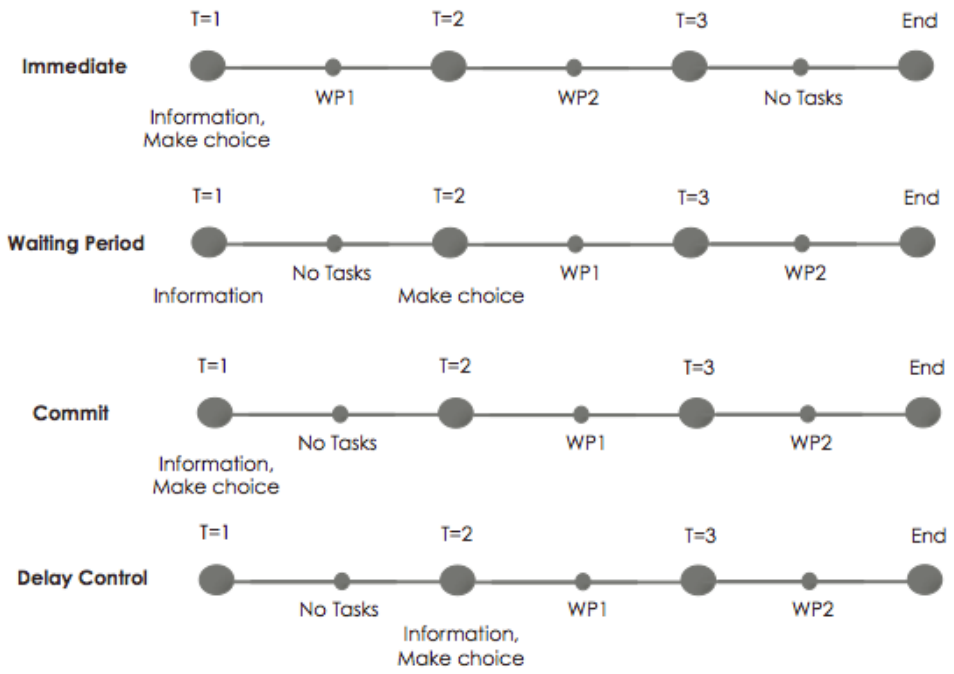

Figure 1: Outline of Experimental Conditions

degree of attention paid to small variations in the interest rate around a critical price point.

To ensure that completing tasks sooner did not result in an earlier end to the experiment, work periods were constrained to last approximately one hour. This was accomplished by setting the work periods to last 60 minutes minus the number of tasks successfully completed. For example, if the participant was required to complete 60 tasks in a period, that period would consist only of those 60 tasks; if she was required to do 40 tasks, the participant would have 20 minutes left over in the work period to engage in a leisure activity of her choice before advancing. This design ensures that participants' decisions only reflected their preferences for allocating effort within the fixed duration of the study – allocation decisions did not change the total length of the study.

We implemented four different treatments. All treatments presented participants with the same budgets as described above and were divided into three one-hour periods. The main difference between treatments was whether WP1 and WP2 were the first two periods or the last two periods. Figure 1 outlines all four treatments.

In the Immediate treatment, participants were presented with information about the budgets and had the opportunity to make their allocation decisions over WP1 and WP2, which began directly after their choice. After the second WP ended, participants had a one-hour period where no

work was required before filling out a questionnaire — the final hurdle that all participants had to clear before receiving payment. In the Waiting Period treatment, participants were presented with information about the budgets but had a one-hour waiting period before having the opportunity to make the allocation choice. As in the Immediate treatment, WP1 and WP2 followed directly after the decision.

The Commit treatment features the same timing of the work periods as in the Waiting Period treatment, with the key difference that participants made allocation decisions *before* the waiting period. In other words, they made fully committed choices over outcomes that were shifted to the future.

Lastly, we designed the Delay Control treatment as a robustness check to ensure that any variation in behavior between the Immediate, Waiting Period, and Commit treatments was due to the waiting period rather than differences in the general timing of work periods. In the Delay Control treatment, participants learned that they would make a decision regarding the allocation of effort tasks but were not presented with the budgets until after a one hour period with no work. The Delay Control treatment has the same timing of WPs as the Waiting Period treatment, but participants did not have a waiting period between being presented with the information and the ability to make a choice. This treatment allows us to rule out alternate explanations such as the delay getting participants used to waiting in general, exogenous information shocks, as well as basic differences in preferences over the timing of outcomes.

In no treatment were participants time constrained when making their allocation decisions. Instead, the crucial difference across treatments was whether the *opportunity* to make a choice was preceded by a waiting period. There were also no differences in the amount of information participants received prior to making their choices. Additionally, as discussed in Cohen et al. (2016), intertemporal decisions over non-monetary consumption events such as food and effort involve potential factors other than the individual's level of time discounting. In the case of effort, factors such as the shape of the effort cost function, complementarities between effort and leisure, or negative anticipatory utility from future effort costs (Loewenstein, 1987) may affect the baseline allocation of tasks between the two periods. It is therefore difficult to interpret the allocation in any one treatment as a pure measure of the individual's time preference.[6] However, given that

[6]For example, in reviewing the literature on measuring time preferences, Cohen et al. (2016) show that estimated

all treatments in our experiment involve the same intertemporal allocation decisions between two periods of the same length, these factors should be constant across conditions. As such, we focus on the differences in choices across treatments as measures of differences in time preferences.

2.1.2 Hypotheses

In this section, we outline the hypotheses that different models of time preferences make in our setting. Consider a decision-maker (DM) who chooses to allocate tasks, x_t, between work periods which can occur in $t = 0, 1, 2$. The DM evaluates utility as $U_k(x_0, x_1, x_2) = \sum_{t=k}^{2} D(t-k)u(x_t)$, where $D(\cdot)$ is the DM's discounting function, and k represents the time period in which the evaluation is made. Assuming that the tasks are not enjoyable to perform, normalizing the instantaneous disutility of effort function $u(0) = 0$ and the discounting function $D(0) = 1$, the DM in our Immediate treatment solves the following decision problem

$$\min_{x_0,x_1} U_0(x_0, x_1) = u(x_0) + D(1)u(x_1) \quad \text{s.t.} \quad x_0 + \frac{x_1}{1+r} = 40 \quad , \tag{1}$$

where r is the interest rate by which tasks avoided in the earlier period grow. In the Waiting Period treatment the allocation choice is shifted by one period, so the DM solves

$$\min_{x_1,x_2} U_1(x_1, x_2) = u(x_1) + D(1)u(x_2) \quad \text{s.t.} \quad x_1 + \frac{x_2}{1+r} = 40 \quad . \tag{2}$$

We first outline the predictions of the standard exponential discounting model and the commonly used behavioral model of present bias with quasi-hyperbolic discounting, which relaxes the assumption of constant discounting. Under constant, exponential discounting, $D(t-k) = \delta^{t-k}$, with δ typically $\in [0, 1]$. The DM in the Waiting Period treatment solves the same decision problem subject to the same constraint as in the Immediate treatment, with the labels shifted by one period. In turn, under exponential discounting the allocations should be the same in both treatments.

Under quasi-hyperbolic discounting, $D(t-k) = \beta^{1(t>k)}\delta^{t-k}$. The β parameter serves to further discount any utility or disutility not received immediately. $\beta \in [0, 1)$ corresponds to a violation of constant discounting and is used to model impulsivity and procrastination (Laibson, 1997). When

discount rates for consumption rewards are consistently higher than for monetary rewards.

the DM first receives information about the decision in $t = 0$ of the Waiting Period treatment, she evaluates the future allocation as

$$\min_{x_1,x_2} U_0(x_1, x_2) = D(1)u(x_1) + D(2)u(x_2) \quad \text{s.t.} \quad x_1 + \frac{x_2}{1+r} = 40 \quad , \tag{3}$$

where $D(1) = \beta\delta$ and $D(2) = \beta\delta^2$. When $k = 0$, her preferred allocation may indeed be different than her choice in the Immediate treatment. However, absent an ability to commit to that preferred allocation, the DM again faces the decision problem represented in (2) after the waiting period elapses — that is, the decision faced in the Immediate treatment shifted by a period. Note that the same logic holds for the case of true hyperbolic discounting, where the discount factor is non-constant between any two periods. In turn, absent the ability to commit, both exponential and quasi-hyperbolic discounting models predict that a waiting period should not affect the allocation decision.

Hypothesis 1. *Exponential/Quasi-hyperbolic/Hyperbolic discounting: absent the ability to commit,* $x_1^{WP} = x_0^I$.

If the DM in the Waiting Period treatment could commit to an allocation in $k = 0$, then choices in that treatment may differ from those in the Immediate treatment. Compare the problems solved by a quasi-hyperbolic DM in equations (1) and (3). When $\beta \in [0, 1)$, $D(1) < \frac{D(2)}{D(1)}$ and the DM displays present bias and prefers a more patient allocation in equation (3) than (1). O'Donoghue and Rabin (1999) show that a DM who is sophisticated about her present bias may take steps to commit herself to following through on the allocation preferred in $k = 0$. Though our experiments were structured to minimize the availability of external commitment devices, participants in the Waiting Period treatment may have been able to mentally commit to a choice in $k = 0$ and follow through on this initial plan in $k = 1$. The Commit treatment makes this commitment opportunity explicit – it features exactly the choice in equation (3) – to test whether the effects of waiting periods are attributable to this commitment mechanism.[7] Any lapse in mental commitment would

[7]Identification of present-bias depends critically on how the 'present' is defined. McClure et al. (2007) demonstrate that delaying the earliest reward by a ten-minute window is sufficient shift it to the future: participants exhibited significant present bias when choosing between rewards at 0,10 and 20 minutes, but exhibited no detectable present bias when all rewards were shifted by ten minutes (with the earliest reward available ten minutes later). Based on this evidence, we believe the one hour period is sufficient to identify present bias (also see also see Balakrishnan et al. (2017) for recent discussion of this topic.

lead the DM in the Waiting Period treatment to make a less patient choice than in the Commit treatment. In turn, if non-constant discounting such as present bias was driving the difference between the Waiting Period and Immediate treatments, then task allocations to the earlier work period in the Commit treatment should be greater than or equal to those in the Waiting Period treatment.

Hypothesis 2. *Quasi-hyperbolic/Hyperbolic discounting: with a perfect ability to commit in the Waiting Period treatment,* $x_1^{WP} = x_1^C$*. With imperfect ability to commit,* $x_1^{WP} < x_1^C$*.*

One potential mechanism for a greater allocation of tasks to the earlier work period is a general shift to more deliberative processing in decision-making. In the basic dual systems framework, System 1 is associated with 'fast,' automatic processing which leads to myopic choices, while System 2 is seen as more deliberative and prudent (Kahneman, 2011). If waiting periods prompt the DM to rely on System 2 as a general tendency, then they should affect decision-making across choices that may or may not have been considered during the delay; waiting periods should decrease reliance on heuristics on choices that were not presented to the DM beforehand.

Building on work in psychology (Gilbert and Wilson, 2007; Wheeler et al., 1997) and neuroscience (Schacter et al., 2007), Gabaix and Laibson (2017) develop a model of 'prospection,' where the DM is uncertain about the true realization of future utility and generates forecasts by deliberating and mentally simulating these events. When presented with a choice, a decision without deliberation is based on a noisy prior belief about utility from future outcomes. Such 'fast' decisions that rely solely one's priors are analogous to the use of heuristics in intertemporal choice (Kahneman, 2003). Deliberation and mental simulations of future outcomes generate noisy unbiased signals of value. These signals are combined with the DM's prior to generate Bayesian posteriors of future utility. Under the assumption that the variance of simulation noise increases with the horizon – intuitively, events that are further away are more difficult to represent mentally – the DM exhibits *as-if* discounting.[8]

Imperfect forecasting of this sort leads the DM to appear more impatient when making 'fast' choices based on her prior than after deliberation and successive simulations of the decision prob-

[8]In the model, the DM's forecasted utility from future events is a function of the 'true' valuation of the event in the absence of noise, $u(\cdot)$, her prior, and the simulation noise. Gabaix and Laibson (2017) show that the two latter components can be captured by an as-if discounting function $D(\cdot)$ that is inversely proportional to the variance of the DM's simulation noise, such that the decision problem can be represented similarly to equations 1-3.

lem. As such, the model can be thought of as providing a micro-foundation for the shift from 'fast' automatic to 'slow,' more deliberative processing in intertemporal choice. Under the assumption that waiting periods prompt deliberation, which leads the DM to generate more mental simulations of future events, her choices should become more patient and less myopic. Conditional on the DM preferring to minimize the number of tasks she needs to complete, all else equal, she will allocate more tasks to the earlier period with a waiting period than without one.[9] Importantly, this result depends on the existence of a tradeoff associated with allocating tasks to the second work period, such that in the absence of simulation noise, allocating tasks to the sooner period is preferred to allocating to the later period. The framework does not predict a difference between treatments when there is no tradeoff — when the interest rate $r = 0$.

Both a general shift toward more deliberative processing and model of 'prospection' predict a difference in allocations after a waiting period:

Hypothesis 3. *Waiting Period:* $x_1^{WP} \geq x_0^I$ *and* $x_1^{WP} \geq x_1^C$.

The 'prospection' framework of Gabaix and Laibson (2017) makes additional predictions. The DM generates simulations of future utility based on information about the choice set presented before the waiting period. Without that information, the waiting period should have no systematic effect on decision-making. Our Delay Control treatment can be used to test this prediction versus a general tendency to allocate tasks to sooner periods, i.e. $x_1^{WP} = x_1^{DC}$ and $x_0^I < x_1^{DC}$. Moreover, the effects of the waiting period should be domain-specific to the choices being considered — behavior on other choices prone to heuristics should not be affected.

Hypothesis 4. *Domain Specificity:* $x_1^{WP} \geq x_1^{DC}$ *and* $x_0^I = x_1^{DC}$. *Additionally, the waiting period will not impact quality of choices unrelated to the allocation choices being considered.*[10]

The framework of Gabaix and Laibson (2017) makes an additional prediction on the relationship between a DM's information processing capacity and her response to waiting periods. Specifically, introducing a waiting period should lead to a larger allocation of tasks to earlier periods for those with greater information processing capacities. We derive this prediction formally in Appendix A.1.

[9] We derive this result formally in Appendix Section A.1.

[10] In the terminology of Halevy (2015), a difference between x_1^{WP} and x_1^{DC} represents a violation of time invariance, which stipulates that temporal allocations should be evaluated relative to "stopwatch time."

In the following section, we present results from the online labor market study testing the hypotheses. We then describe the setup of the laboratory experiment which provides a further test of the predictions.

2.1.3 Results

Participants' allocation decisions were responsive to the interest rates. Examining decisions on the convex budgets, participants allocated significantly more tasks to WP1 as the interest rate increased in all four treatments.[11] Considering the binary choices, the majority of subjects allocated all tasks to Work Period 2 when the interest rate was negative while the majority allocated all tasks to Work Period 1 when the interest rate was marginally positive (difference $= 46\%, p < 0.001$). Together, these results confirm that participants were attentive to our manipulation of the interest rate. Only one of 122 subjects switched against the grain on the binary choices.

Turning to treatment differences, we examine allocation decisions on the convex budgets. Comparing the Immediate and Waiting Period treatments allows us to test Hypothesis 1 versus Hypothesis 3. As illustrated in Figure 2, Panel A, the results support Hypothesis 3. Participants allocated significantly more tasks to WP1 in the Waiting Period treatment than in the Immediate treatment ($p = 0.02$).[12] The magnitude of this effect is large: about 17% of the Immediate mean, or half of a standard deviation.

We now examine the effect of the waiting period by budget. Budgets 1-3 offer a tradeoff between completing fewer tasks sooner versus a larger amount later, with implied interest rates of 50%, 25% and 12.5%, respectively. Decisions on Budget 4, which had a 0% implied interest rate, did not involve a tradeoff between fewer tasks now versus more tasks later. Regression estimates of the treatment effect on tasks allocated to WP1, separately by budget, are presented in Table 2.[13] Participants in the Waiting Period treatment allocated significantly more tasks to WP1 than those in the Immediate treatment across all three budgets with positive interest rates. The effect sizes are

[11] Estimates are from random effects models of tasks allocated to WP1 as a function of $100 \cdot ln(1+r)$. Standard errors are clustered at the individual level. $\beta = 0.09$ ($p < 0.01$), 0.12 ($p = 0.01$), 0.23 ($p < 0.01$), and 0.11 ($p = 0.01$) for Immediate, Waiting Period, Delay Control and Commit, respectively.

[12] Estimate is from a random effects model of tasks allocated to WP1 as a function of treatment dummy variables. Standard errors are clustered at the individual level.

[13] Simultaneous specifications that allow for individual errors to correlate across the four estimating equations yield almost identical results.

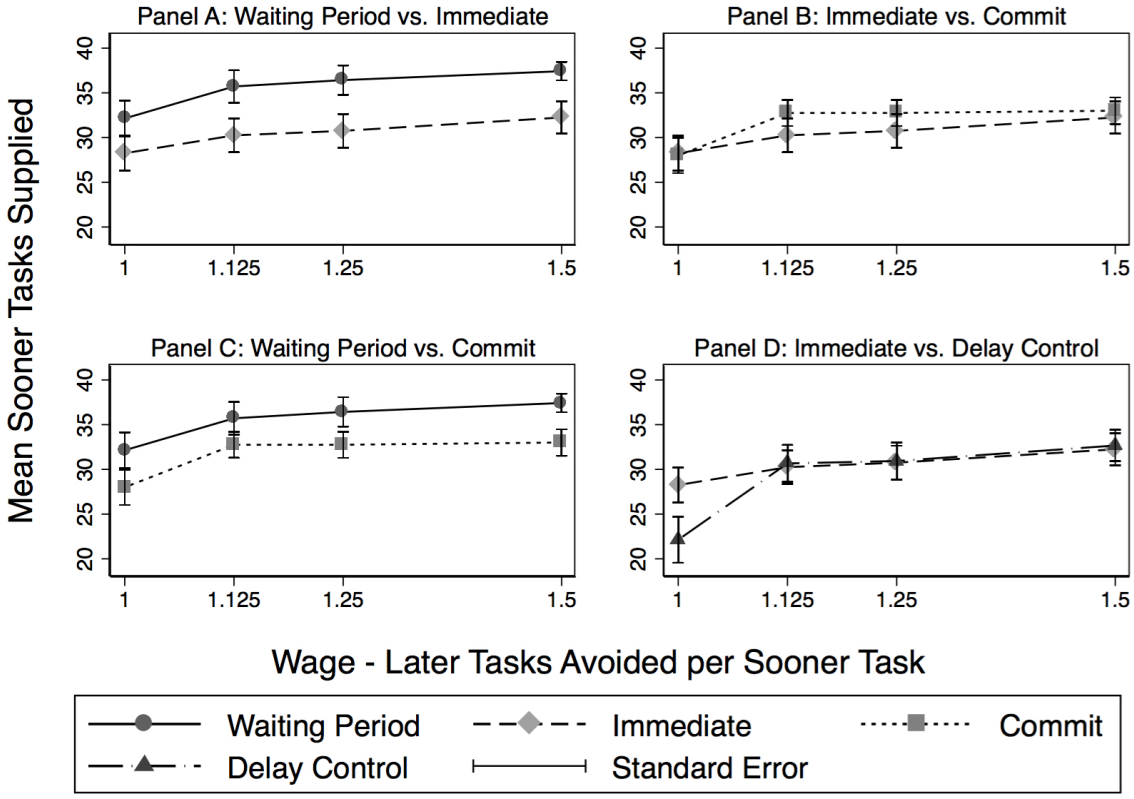

Figure 2: Tasks Performed in Work Period 1, by Treatment

about half of a standard deviation in each instance.

Importantly, waiting periods did not simply prompt participants to allocate tasks to the earlier period in general. Waiting periods only led to significantly earlier allocations if this resulted in fewer tasks to complete overall – on budgets with positive interest rates. The treatment effect on the 0% interest rate budget shrinks by 28% and is not statistically different from zero.[14] Considering the binary-choice budget with a negative interest rate –meaning that subjects could do fewer tasks by waiting until WP2– we find that waiting periods made subjects 13% *more* likely to allocate all tasks to WP2 (though the effect is not statistically significant, $p = 0.29$, marginal effect from probit model). When the interest rate is positive, waiting periods lead to more tasks being allocated to the sooner period; when the interest rate is negative, tasks are (directionally) more likely to be allocated to the later period. Together, these results offer suggestive evidence for individuals becoming better calibrated after waiting periods, rather than just shifting tasks to the later period in general.[15]

[14]It is not statistically different from the estimates of the waiting period effect on the other budgets, however.

[15]Consistent with this, on the binary choice budget with an interest rate of essentially zero (40 tasks in WP1 vs. 41

Table 2: Effect of Treatment on Convex Task Allocations to Work Period 1

Interest rate:	50%	25%	12.5%	0%
	(1)	(2)	(3)	(4)
Waiting Period	5.179**	5.679**	5.464**	3.893
	(2.029)	(2.446)	(2.562)	(2.766)
Commit	0.750	2.000	2.500	-0.250
	(2.234)	(2.340)	(2.350)	(2.702)
Delay Control	0.417	0.183	0.417	-6.117*
	(2.467)	(2.813)	(2.850)	(3.193)
Constant	32.250	30.750	30.250	28.250
	(1.757)	(1.869)	(1.899)	(1.927)
N	122	122	122	122

** : $p < 0.05$, * : $p < 0.10$. Coefficients from OLS models. Bootstrapped standard errors from 1000 replications are reported in parentheses below each estimate. Output is reproducible with a seed of 1.

Figure A1 in the appendix shows the full distribution of choices on the positive-interest convex choice sets. Because of the frequency of subjects choosing to do all their tasks in WP1, we use a Tobit model to account for corner solutions and separately estimate the extensive and intensive margin in Appendix Table A1. We find that the waiting period affects both the probability of doing all tasks in WP1 and the number of tasks done, conditional on not selecting a corner solution. Together, these results suggest that introducing a waiting period between information about a choice and the choice itself leads to more patient decisions. As such, we can reject Hypothesis 1 in favor of Hypothesis 3.

Comparing decisions in the Commit treatment to those in the Waiting Period treatment allows us to test Hypothesis 2: whether the effects of waiting periods operate via sophisticated present bias and the ability to mentally commit to a plan. As shown in Figure 2, Panel C, participants in the Waiting Period treatment allocate significantly more tasks to WP1 than those in the Commit treatment across the budgets ($p = 0.04$). This suggests that the effect of the waiting period cannot be explained solely by models that relax the assumption of constant discounting such as present bias. Table 2 shows the effects of the Commit and Waiting Period treatment relative to the Immediate treatment by budget. Comparing the Waiting Period coefficients to the Commit coefficients,

tasks in WP2) we again find no impact of the waiting period (marginal effect of -0.01, $p = 0.89$).

the Waiting Period coefficient is at least twice as large on all positive interest rate budgets. We can reject equality of the coefficients on budgets with $r = 0.5$ and $r = 0.25$ ($p = 0.01$ and $p = 0.08$, respectively). We fail to reject equality when $r = 0.125$ ($p = 0.19$), although the difference remains sizable.[16] These results are robust to the Tobit model specification in Appendix Table A1.

As a robustness check and test of Hypothesis 4, we compare the Waiting Period treatment to choices in Delay Control. Participants in the former treatment allocated significantly more tasks to WP1 than in the latter, on average, across the convex budgets ($p < 0.01$). In Table 2, the coefficient on the Waiting Period treatment is significantly different from the Delay Control treatment across all budgets.[17] For the positive-interest rate budgets, allocation decisions in the Delay Control treatment were not significantly different from those in the Immediate treatment, as shown in Figure 2, Panel D. Participants in the Delay Control treatment did seem to allocate *fewer* effort tasks to WP1 on Budget 4 (0% interest rate) compared to those in the Immediate treatment, though this difference was only marginally significant. In line with Hypothesis 4, waiting periods appear to only affect decisions that can be considered based on information available beforehand, which is not consistent with a general shift to more deliberative decision-making.

In Appendix Section A.2, we estimate the parameters of an intertemporal utility function that permits both present bias and an as-if discounting simulation parameter. Structural estimation of such parameters is common in the experimental discounting literation. This allows us to compare our estimate of present bias over effort to other estimates in order to gauge the comparability of our procedure.[18] We identify present bias of similar magnitude to estimates in prior work on effort: $\beta = 0.912$ (S.E. = 0.042, $p = 0.037$, tested against $\beta = 1$).[19] We also estimate a "simulation parameter," S, that can be interpreted as a multiplicative factor that scales the discount factor when there is no time for deliberation. Like β, S is significantly less than one: $S = 0.841$ (S.E. = 0.046, $p = 0.001$, tested against $S = 1$). S is also significantly different from β ($p = 0.045$). In

[16] For Budgets 1-3, the coefficients of the Commit treatment are all positive though not significant, providing suggestive evidence of present bias. This is also shown in Figure 2, Panel B. In the Appendix we estimate structural parameters that identifies a present bias parameter $\beta < 1$ in a model of quasi-hyperbolic discounting. However, consistent with the reduced form results, it does not explain the full impact of waiting periods on choices.

[17] For $r = 0.5$, $p = 0.02$, for $r = 0.25$, $p = 0.04$, for $r = 0.125$, $p = 0.06$, and for $r = 0$, $p < 0.01$.

[18] As discussed earlier in the section, the short-horizon nature of the study and factors orthogonal to time preferences in participants' effort allocation decisions complicate inference of long-horizon discount rates, δ. In turn, our estimation procedure focuses on short-horizon parameters that can be identified directly through treatment variation.

[19] Augenblick et al. (2015) estimate $\beta = 0.90$.

sum, while our design captures significant present-bias, the independent effect of waiting periods on intertemporal decisions is statistically larger. Details on the utility maximization problem, its solution, the censored-Tobit maximum likelihood procedure, and the full set of estimates are in the Appendix.

2.2 Laboratory Study

2.2.1 Design and Implementation

In the online labor market study, we lacked precise information on what participants did during the waiting periods. Though end-line survey responses suggest that participants largely took this time to engage in unrelated activities like reading a book or cooking a meal, the setup does not allow us to completely rule out alternative mechanisms. For example, participants may have used the waiting period to consult with others.

To investigate whether such alternative explanations could be driving the waiting period effect, we replicated the effort study in the laboratory where activities available to the participants were limited and observable. We recruited participants (N = 72) for a study at Carnegie Mellon University's Center for Behavioral Decision Research. The advertised study duration was 3.5 hours, with a reward of $50 for completion of the study.

Participants in the study faced the same choices and the same interface as those in the online study. Upon arriving to the lab, participants were randomized into either the Immediate or the Waiting Period treatment. One major difference in protocol for the lab study was that we were able to explicitly prohibit communication and limit the types of activities participants could engage in during the waiting period and other intervals of free time. No cell phone usage was allowed, and internet use was limited to watching media (access to an online streaming site and headphones were provided). Subjects were also encouraged to read material they had brought to the lab. Research assistants were instructed to monitor activity during waiting periods and free time, and responses to end line surveys are consistent with participants engaging in the suggested activities.[20]

We collected two additional types of data in this study to allow for further tests of the hypotheses. First, we embedded questions designed to test the propensity for automatic decisions versus

[20]One participant was observed trying to use a cell phone, and is excluded from the analysis. Inclusion of their data does not affect our results.

deliberative processing immediately following the task allocation choices in both treatments; in one treatment participants had a waiting period prior to answering these questions, in the other they did not. The set consisted of seven Cognitive Reflection Test questions (CRT, Frederick (2005), CRT2 Thomson and Oppenheimer (2016)).[21] Participants were not given information about these questions at the onset of the experiment; unlike the task allocation decision, participants in the Waiting Period treatment could not consider them during the waiting period. If waiting periods lead to more deliberative processing in general, then we should observe a treatment effect on the number of questions answered correctly. However, because these questions are in a different domain than the information presented before the waiting period, Hypothesis 4 predicts that there should be no such effect.

We also included four belief-bias questions from Markovits and Nantel (1989) which have been verified by Thomson and Oppenheimer (2016) as a measure of cognitive sophistication and information processing capacity.[22] These questions were asked during the end-line survey presented to participants at the same time across both treatments. Correct answers on two of the questions required accepting the conclusion; correct answers on the other two questions required rejecting the conclusion. Since the aim was to detect variation in cognitive sophistication and information processing, we created a measure that would in expectation distinguish between those who randomly guessed or accepted/rejected every conclusion, and those who correctly evaluated the logical validity of the arguments. The three-level variable differentiated those who got two or fewer questions correct, three questions correct, or all four correct.[23] A larger impact of waiting periods for individuals who scored higher on this measure would provide further evidence for theory based on 'prospection.'

[21] For example, participants were asked "A bat and a ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost?" The 'fast' automatic response is $0.10; overiding this heuristic response through deliberative processing leads to the correct answer, $0.05.

[22] Belief bias is defined as the propensity to agree with a conclusion that seems believable when judging the logical validity of an argument, but does not follow from the evidence provided. Participants were given two statements and a conclusion, and asked whether the statements were sufficient to derive the conclusion. For example, "Premise 1: All things that are smoked are good for the health; Premise 2: Cigarettes are smoked; Conclusion: Cigarettes are good for the health."

[23] The results are robust to alternative specifications of the variable, such as having it correspond to the total number correct.

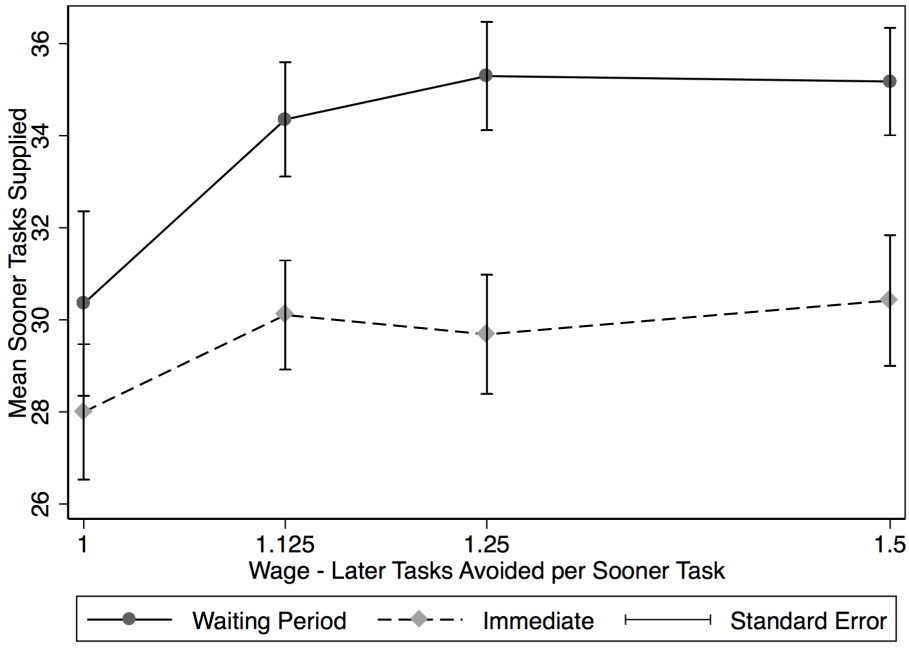

Figure 3: Tasks Performed in Work Period 1, by Treatment, Lab Study

2.2.2 Results

Results from the laboratory are consistent with those from the online labor market. The average number of tasks done in WP1 across all convex tasks is 31.6 ($S.D. = 8.7$) in the lab study and 31.5 ($S.D. = 10.6$) in the online labor market study. Participants were similarly responsive to changes in the interest rate.[24]

Consistent with the third hypothesis, participants allocated significantly more tasks to WP1 in the Waiting Period treatment than in the Immediate treatment ($p = 0.01$).[25] This is shown in Figure 3. The 14% effect size is very similar to our estimate of 17% in the online labor market study. In both studies, this is roughly half a standard deviation.

Table 3 presents results separately by budget. Again, participants in the Waiting Period treatment allocated significantly more tasks to WP1 than did those in the Immediate treatment across all three convex budgets with a positive interest rate.[26] The size of the effect shrinks by 51%, moving

[24] Using a random effects models of tasks allocated to WP1 as a function of $100 \cdot ln(1 + r)$ with standard errors are clustered at the individual level: $\beta = 0.08$ ($p < 0.01$).

[25] Estimate is from a random effects model of tasks allocated to WP1 as a function of a Waiting Period dummy variable. Standard errors are clustered at the individual level.

[26] As with the online labor market study, we bootstrap standard errors for each budget, and they are very similar using a simultaneous estimation with clustered standard errors.

Table 3: Effect of Treatment on Convex Task Allocations to Work Period 1
Lab Study

Interest rate:	50%	25%	12.5%	0%
	(1)	(2)	(3)	(4)
Waiting Period	4.755***	5.610***	4.248**	2.353
	(1.791)	(1.721)	(1.712)	(2.416)
Constant	30.421	29.684	30.105	28.000
	(1.406)	(1.288)	(1.204)	(1.429)
N	72	72	72	72

** : $p < 0.05$, * : $p < 0.10$. Coefficients from OLS models. Bootstrapped standard errors from 1000 replications are reported in parentheses below each estimate. Output is reproducible with a seed of 1.

from the positive-interest budgets to the zero-interest budget, and is no longer significantly different from zero. Figure A2 in the appendix shows the full distribution of choices on the positive-interest convex choices. Results are qualitatively similar for the Tobit estimates shown in Appendix Table A2. Overall, these results are remarkably consistent with those from the online labor market study, indicating that the observed effects were not driven by access to communication or other activities pursued during the waiting period.

Next, we consider the impact of the waiting period on choices unrelated to the intertemporal task budgets. If the waiting period leads to a general shift towards more deliberative decision-making, we expect to see higher scores on the Cognitive Reflection Task in that treatment. This is not what we find. In the Immediate treatment, the average CRT score is 1.98 correct answers (out of three), and in the Waiting Period treatment, the average score is 2.01 (difference: $p = 0.92$). The average CRT2 score is 2.52 in Immediate and 2.48 in Waiting Period (difference: $p = 0.93$). Consistent with Hypothesis 4, the waiting period appears to have had little to no impact on decision quality outside the domain of the task allocation.

Finally, we examine the relationship between information processing capacity and response to the waiting period. First examining choices in the Waiting Period treatment only, we regress the number of tasks allocated to WP1 on our measure of information processing capacity. Pooling all the positive-interest convex choices we estimate a random effects model with standard errors clustered at the individual level. Because information processing capacity may be correlated with other individual traits that determine preferences for allocating tasks between now and later, we also ex-

Table 4: Effect of Treatment on Positive-Interest Convex Task Allocations to Work Period 1 By Belief Bias Level

Sample:	Waiting Period		All	
	(1)	(2)	(3)	(4)
Information Processing	2.447	3.028**	-0.631	-1.168
	(1.635)	(1.284)	(1.351)	(0.963)
$r = 0\%$		0.310**		0.399***
		(0.123)		(0.093)
Waiting Period			1.528	-0.805
			(2.915)	(1.651)
Waiting Period X Information Processing			3.077	4.362***
			(2.114)	(1.564)
Constant	32.278	22.226	30.751	20.148
	(2.349)	(4.000)	(1.745)	(2.727)
Clusters	34	34	72	72
N	102	102	216	216

*** : $p < 0.01$, ** : $p < 0.05$. Coefficients from Random Effects models. Standard errors clustered by individual are reported in parentheses below each estimate. Information processing ability is low, medium or high depend on whether subjects answered two or fewer, three, or four belief bias question correctly.

amine the relationship conditional on participants' choices on the zero-interest budget, which did not present a tradeoff for intertemporal allocations. These results are presented in columns (1) and (2) of Table 4. The relationship between information processing capacity and allocating tasks to the earlier period is positive, but only statistically significant when controlling for choices on the zero-interest budget.[27]

Looking across treatments, we regress allocated tasks to WP1 on a treatment dummy variable, our measure of information processing capacity, and their interaction. As before, we also control for decisions on the $r = 0\%$ budget.[28] Results are presented in columns (3) and (4) of Table 4. In column (3), going from lowest to highest level of information processing leads to a three-task increase in the size of the treatment effect, although this estimate is not statistically significant.

[27] Results using a Tobit model are similar, presented in Appendix Table A3.

[28] While this is potentially an endogenous control variable, the induced bias means that our estimates of the treatment effect in these specifications should be a lower bound. The waiting period has a positive, but not statistically significant impact on Work Period 1 tasks on the zero-interest budget. Therefore, holding those tasks constant, we would expect an individual in Immediate to be more patient than an individual in Waiting Period. This would bias the estimates of the treatment on Work Period 1 tasks downwards.

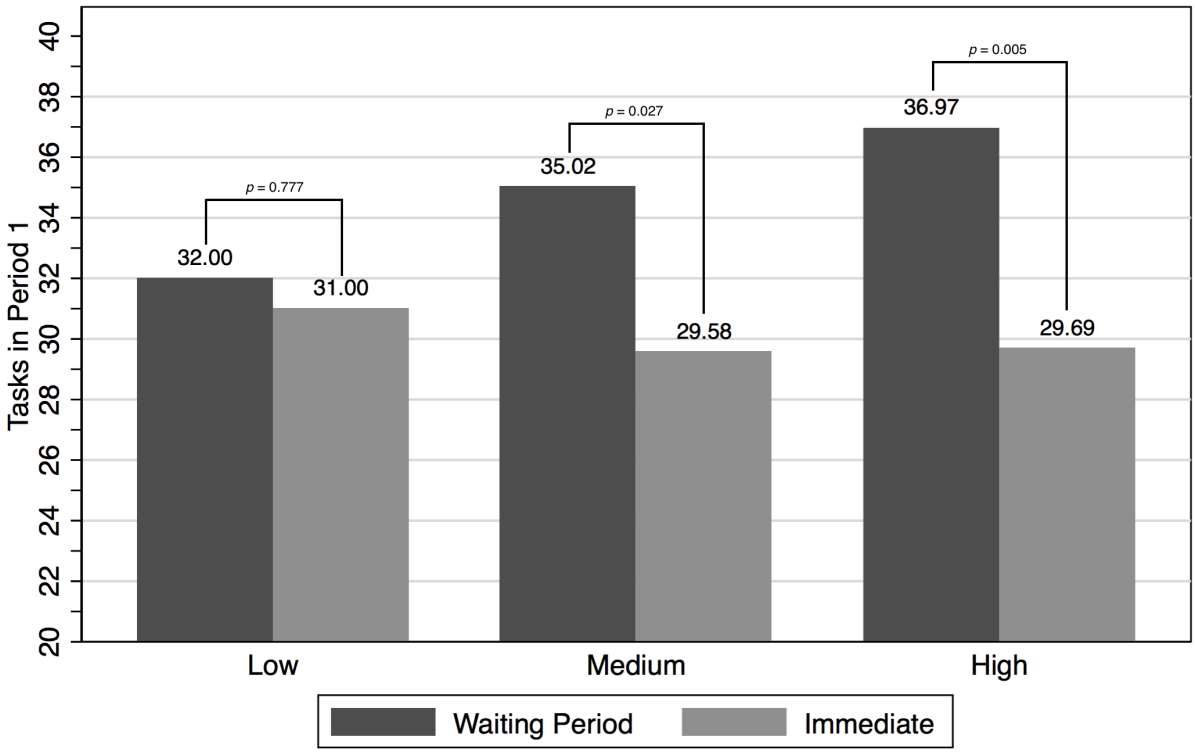

Figure 4: Positive-interest Convex Task Allocations by Information Processing Ability and Treatment

In column (4), where we control for the $r = 0\%$ task allocation, the treatment effect increases with processing capacity by more than four tasks allocated to the earlier work period. This effect is statistically significant at a 1% level. Figure 3 shows the treatment effects separately by our information processing classification.[29] Though this evidence is supportive of the 'prospection' account, it should be taken as suggestive – though the effects are consistently positive, they are significant only for some specifications.

3 Waiting Periods for Consumption Goods

3.1 Design and Implementation

Our third study demonstrates the application of waiting periods in a field context. We partnered with a small grocery store in a residential area in Bukavu, a city on the Eastern border of the

[29] p-values from Random Effects models with standard errors clustered by individual.

Democratic Republic Congo (DRC). The store sells everyday goods and simple foodstuffs like rice, water, and milk. It also has access to electricity and refrigeration that is lacking in most homes, and the vast majority of the people in our sample visited the store every day to pick up groceries. The store ran as usual during the study and was staffed by the family that has owned and operated it for the past decade in order to avoid disrupting customers' familiarity with the store and to reduce uncertainty related to the experiment taking place. One of the authors supervised all aspects of the procedures for the entire length of the experiment.

A total of 258 store customers participated in the study, where each made a decision of when to redeem a coupon for a set amount of flour. Cassava flour is a staple crop and consumption good in the Eastern DRC, particularly for making *fufu* and *chikwangue*, dough preparations. Cassava products, in general, contribute to about 65% of daily calories consumed in the DRC, and are the main food crop for 80% of the population, which translates to about 0.4kg per-capita daily – mostly from fufu and chikwangue (Harvest Plus, 2010). The coupons were were redeemable for flour that the store typically sold, not a new product.

Upon arriving at the store and agreeing to participate, all customers completed a demographic survey.[30] Participants who were illiterate or had difficulty completing the survey on their own were helped by a research assistant who was blind to the hypothesis. The survey was in both Swahili and French and the participant chose which was more convenient for them. On average the survey took 30 minutes to complete.[31]

Participants were then randomly assigned to one of two treatments — Immediate or Waiting Period. In both, they received a coupon that could be exchanged for varying amounts of flour depending on the day it was redeemed.[32] In the Immediate treatment, the participant could redeem the coupon on the same day for 1 bag of flour (approximately 1kg). If she chose not to redeem it on that day, she could come back the next day for 2 bags, and so on, up until 5 bags of flour. The Waiting Period treatment shifted the redemption schedule by one day: the participant had to wait a day before deciding whether to redeem the coupon for 1 bag of flour. As in the Immediate treatment, if she chose not to redeem the coupon on the day after the waiting period, she could

[30]The survey was presented in the beginning of the experiment in order to collect demographic information and other variables of interest in case of differential attrition. As discussed further below, attrition turned out to be minimal.
[31]Full questionnaire available upon request.
[32]Each coupon had an ID matching it with a questionnaire, a date of issue and a code signifying the treatment.

Table 5: Coupon Value over Time - kg of Flour

Treatment:	Immediate	Waiting Period
Day of Receipt	1	0
1 Day after Receipt	2	1
2 Days after Receipt	3	2
3 Days after Receipt	4	3
4 Days after Receipt	5	4
5 Days after Receipt	0	5
6 Days after Receipt	0	0

come back the next day for two bags of flour, and so on, up until 5 bags of flour. Table 5 presents the coupon value schedule. The study was run over a number of days, meaning that treatment status and the calendar date of redemption options are not co-linear. Due to the material incentives and participants' daily visits to the store, only one participant did not redeem their coupon by the last possible day (this individual was in the Waiting Period treatment).

There may be transaction costs associated with returning to the store, which are constant across all redemption dates in the Waiting Period treatment but not in the Immediate treatment because participants were recruited at the store. We note that participants coming to the store to shop presumably had means to purchase food without the existence of our study, and that the store was a common daily destination for consumers. In addition, we collect data on both food access and the distance participants live from the store. We also use data on risk and trust attitudes as controls for other factors that may have differentially influenced choices in the Immediate and Waiting Period treatments. Responses on the questionnaire are used to verify that key demographic and preference variables were uncorrelated with treatment assignment. The frequency of significant differences is consistent with random assignment (see Appendix Table A3). Most importantly, neither measures of trust of others, stated preference for risk, nor food access were correlated with treatment assignment.

Our key dependent measure is the likelihood that an individual redeemed her coupon for its minimum value of one bag of flour. Once an individual in the Immediate treatment chooses not to redeem their coupon on the day it is received, she experiences an overnight waiting period.

Additionally, given the expected difference in minimum-value redemption rates, the individuals who chose not to redeem on the earliest possible date in the Immediate treatment will be more selected relative to those who chose not to redeem in the Waiting Period treatment. Particularly, those who resist the urge to redeem in the Immediate treatment may be more patient on average than those who resist after a waiting period, making comparisons of choices after the earliest possible redemption date subject to issues of selection. As a result, studying minimum-value redemption provides the closest analog to our online and lab studies.

3.2 Results

The introduction of a waiting period has a substantial effect on minimum-value redemption rates: 34 individuals (25%) in the Immediate treatment redeemed the coupon on the earliest possible date, compared to 11 (9%) in the Waiting Period treatment. The 16 percentage-point difference is statistically significant ($p < 0.01$). Results are presented in Table 6. This estimate is robust to the addition of control variables and their interactions with the treatment variable: food access, distance from the store, trust in other and risk tolerance do not diminish the size of the effect. Table A4 in the Appendix presents the corresponding Probit model estimates. The impact of the waiting period is slightly larger in this case.

Examining choices after the earliest redemption date suggests that indeed, the sample of individuals who resisted the urge to redeem the coupon in the Immediate treatment were more patient after an overnight waiting period than those who resisted in the Waiting Period treatment. Conditional on not redeeming earliest possible date, those in the Waiting Period treatment redeemed their coupon 0.73 days sooner than those in the Immediate treatment ($p < 0.01$).

Despite attempts to minimize potential confounds such as differential transaction costs, the difficulties endemic to running the study in the field prevent us from completely ruling them out. In this sense, the findings from this study are complementary to our lab and online studies where such concerns are minimized. Observing analogous effects of waiting periods on intertemporal decisions across all three settings points to the generalizability of the results and the applicability of waiting periods as a potential intervention.

Table 6: Impact of Waiting Period on Likelihood of Minimum-value Coupon Redemption

	(1)	(2)
Waiting Period	-0.160***	-0.166***
	(0.045)	(0.047)
Food Access		-0.020
		(0.044)
Food Access X Waiting Period		0.013
		(0.050)
Distance from Store		0.037
		(0.060)
Distance from Store X Waiting Period		$>$ -0.001
		(0.069)
Trust in Others		-0.003
		(0.041
Trust in Others X Waiting Period		0.014
		(0.048)
Risk Tolerance		-0.020
		(0.044)
Risk Tolerance X Waiting Period		0.067
		(0.050)
Constant	0.250	0.253
	(0.037)	(0.039)
N	258	252

*** : $p < 0.001$. Coefficients from OLS models. Robust standard errors are reported in parentheses below each estimate. Food access, trust in others and risk tolerance are measured on 1-4 scales. Distance from the store is measured on a 1-3 scale. All control variables are de-meaned. We lose six observations with the addition of control variables due to incomplete survey responses.

4 Conclusion

Across three studies, we demonstrate the significant effect of prompting deliberation through waiting periods on intertemporal choices. When an intertemporal choice is preceded by a waiting period, participants in an effort allocation task choose to complete more tasks earlier, thereby minimizing overall work time and unpleasant effort. This effect cannot be explained by non-constant discounting such as present-biased time preferences. In our first study, the effect of temporal sepa-

ration between the receipt of information about a choice and the ability to make it is stronger than the effect of temporal separation between a choice and its consequences, i.e. the ability to commit to a future allocation. These findings are consistent with theory that explicitly considers the role of deliberation in intertemporal choice, such as the framework of Gabaix and Laibson (2017). Lastly, we demonstrate an application of waiting periods in a field setting, showing that grocery store customers are less likely to redeem their coupon for its minimum value of flour when the initial choice of whether to redeem or not is preceded by a waiting period.

Our results have implications for policy aimed at affecting people's intertemporal choices. Economists have noted the lack of demand for commitment devices, which contrasts with the predictions of discounting models designed to capture myopic behavior (Laibson, 2015).[33] Our findings suggest that the difference in choices over future periods versus choices between the present and the future – the hypothesized driver of demand for commitment – is relatively small compared to choices with or without waiting periods. Additionally, unlike commitment contracts, waiting periods are common in the real world. Encouraging deliberation through the use of waiting periods for intertemporal decisions may be an underutilized policy tool that is not costly and less paternalistic than forced commitment, but may have substantial benefits for choices over time.

Since waiting periods do not restrict individuals' choice or information sets, they may represent a more feasible policy tool for encouraging patience in decision-making than other interventions. Consider the tax refund example from the introduction. Our results suggest that eliminating the waiting period between being informed of the refund and the ability to use the windfall could have a significant impact on the choice of whether to spend or save the money. A firm offering to deliver the refund immediately through an Anticipation Loan, for example, may create substantial negative downstream consequences for the consumer even if the loan were interest-free.

Additional work is needed to determine if these results extend to longer horizons and choices with higher stakes. All choices in our studies involved consumption goods: leisure or flour. Examining the impact of waiting periods on intertemporal money allocations like saving and borrowing would be particularly interesting, especially given the large literature contrasting consumption and money discounting. Another promising avenue for further research would study the effects of wait-

[33] It should be noted that Laibson (2015) shows that demand for commitment is actually predicted by models of quasi-hyperbolic discounting for only a narrow set of parameters, which could also explain its paucity.

ing periods on the intensive margin, examining the length of delay necessary to influence patience.

References

J. Abeler, A. Falk, L. Goette, and D. Huffman. Reference points and effort provision. *American Economic Review*, 101(2):470–492, 2011.

S. Andersen, G.W. Harrison, M. I. Lau, and E. E.Rutstrom. Discounting behavior: A reconsideration. *European Economic Review*, 71:15–33, 2014.

J. Andreoni, M.A. Kuhn, and C. Sprenger. Measuring time preferences: A comparison of experimental methods. *Journal of Economic Behavior and Organization*, 116:451–464, 2015.

N. Augenblick, M. Niederle, and C. Sprenger. Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics*, 130(3):1067–1115, 2015.

Uttara Balakrishnan, Johannes Haushofer, and Pamela Jakiela. How soon is now? evidence of present bias from convex time budget experiments. Working Paper 23558, National Bureau of Economic Research, June 2017. URL http://www.nber.org/papers/w23558.

N. Barberis, A. Mukherjee, and B. Wang. First impressions: "system 1" thinking and the cross-section of stock returns. Working paper, 2013.

J. Benhabib, A. Bisin, and A. Schotter. Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior*, 69(2):205–223, 2010.

A. W. Brooks. Emotion and the art of negotiation. *Harvard Business Review*, 2015.

H. Burgess. Cooling-off periods. In G. Burgess and H. Burgess, editors, *Beyond Intractability*. Conflict Research Consortium, University of Colorado, Boulder, Colorado, 2004.

Andrew Caplin and Mark Dean. Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review*, 105(7): 2183–2203, jul 2015. ISSN 0002-8282. doi: 10.1257/aer.20140117. URL http://pubs.aeaweb.org/doi/10.1257/aer.20140117.

Andrew Caplin, Mark Dean, and Daniel Martin. Search and Satisficing. *American Economic Review*, 101(7):2899–2922, dec 2011. ISSN 0002-8282. doi: 10.1257/aer.101.7.2899. URL http://pubs.aeaweb.org/doi/10.1257/aer.101.7.2899.

Jonathan D. Cohen, Keith Marzilli Ericson, David Laibson, and John Myles White. Measuring time preferences. Working Paper 22455, National Bureau of Economic Research, July 2016. URL http://www.nber.org/papers/w22455.

X. Dai and A. Fishbach. When waiting to choose increases patience. *Organizational Behavior and Human Decision Processes*, 121(2):256–266, 2013.

G. S. Edwards, E. Nesson, J.J. Robinson, and F. E. Vars. Looking down the barrel of a loaded gun: The effect of mandatory handgun purchase delays on homicide and suicide. University of Alabama Legal Studies Research Paper No. 2629397, 2016.

K. M. Ericson, J.M. White, D. Laibson, and J.D. Cohen. Money earlier or later? simple heuristics explain intertemporal choices better than delay discounting does. *Psychological Science*, 26(6): 826–833, 2015.

A. Falk, D. Huffman, and K. Mierendorff. Incentive properties and political acceptability of workfare: Evidence from real effort experiments. Working paper, 2006.

S. Frederick. Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4): 25–42, 2005.

X. Gabaix and D. Laibson. Myopia and discounting. National Bureau of Economic Research Working paper No. 23254, 2017.

D. T. Gilbert and T. D. Wilson. Prospection: Experiencing the future. *Science*, 317(5843):1351–1354, 2007.

Y. Halevy. Time consistency: Stationarity and time invariance. *Econometrica*, 83(1):335–352, 2015.

Harvest Plus. Country crop profile: Provitamin a cassava in the democratic republic of congo, 2010.

D. Kahneman. Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5):1449–1475, 2003.

D. Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.

J. Kessler, H. Kivimaki, and M. Niederle. Thinking fast and slow: Generosity over time. Working Paper, 2017.

K. N. Kirby, N. M. Petry, and W. K. Bickel. Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology: General*, 128(1): 78–87, 1999.

C. Koenig and D. Schindler. Dynamics in gun ownership and crime -evidence from the aftermath of sandy hook. Working paper, 2016.

N. Lacetera, D. G. Pope, and J. Sydnor. Heuristic Thinking and Limited Attention in the Car Market. *American Economic Review*, 102(5):2206–2236, 2012.

D. Laibson. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2): 443–477, 1997.

D. Laibson. Why don't present-biased agents make commitments? *American Economic Review Papers and Proceedings*, 105(5):267–272, 2015.

George Loewenstein. Anticipation and the valuation of delayed consumption. *The Economic Journal*, 97(387):666–684, 1987. ISSN 00130133, 14680297. URL http://www.jstor.org/stable/2232929.

Michael Luca, Deepak Malhotra, and Christopher Poliquin. Handgun waiting periods reduce gun deaths. *Proceedings of the National Academy of Sciences*, 2017. doi: 10.1073/pnas.1619896114. URL http://www.pnas.org/content/early/2017/10/11/1619896114.abstract.

H. Markovits and G. Nantel. The belief-bias effect in the production and evaluation of logical conclusions. *Memory and Cognition*, 17(1):11–17, 1989.

S. M. McClure, K.M. Ericson, D. Laibson, G. Loewenstein, and J.D. Cohen. Time discounting for primary rewards. *Journal of Neuroscience*, 27(21):5796–5804, 2007.

J. Metcalfe and W. Mischel. A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review*, 106(1):3–19, 1999.

M. J. Moore and W. Viscusi. Models for estimating discount rates for long-term health risks using labor market data. *Journal of Risk and Uncertainty*, 3(4):381–401, 1990.

T. O'Donoghue and M. Rabin. Doing it now or later. *American Economic Review*, 89(1):103–124, 1999.

D. Read, S. Frederick, and M. Scholten. Drift: An analysis of outcome framing in intertemporal choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(2):573–588, 2013.

A. Rees-Jones and D. Taubinsky. Heuristic perceptions of the income tax: Evidence and implications for debiasing. working paper, 2016.

A. Rubinstein. "economics and psychology"? the case of hyperbolic discounting. *International Economic Review*, 44:1207–1216, 2003.

A. Rubinstein. Response time and decision making: An experimental study. *Judgment and Decision Making*, 8(5):540–551, 2013.

Daniel L Schacter, Donna Rose Addis, and Randy L Buckner. Remembering the past to imagine the future: the prospective brain. *Nature Reviews Neuroscience*, 8(9):657–661, 2007.

R. H. Thaler. Some empirical evidence on dynamic inconsistency. *Economic Letters*, 8(3):201–207, 1981.

R.H. Thaler and H.M. Shefrin. An economic theory of self-control. *Journal of Political Economy*, 89(2):392–406, 1981.

K. S. Thomson and D.M. Oppenheimer. Investigating an alternate form of the cognitive reflection test. *Judgment and Decision Making*, 11(1):99–113, 2016.

Mark A. Wheeler, Donald T. Stuss, and Endel Tulving. Toward a theory of episodic memory: The frontal lobes and autonoetic consciousness. *Psychological Bulletin*, 121(3):331–354, 1997.

A Appendix for Online Publication

A.1 Relationship between Information Processing Capacity and Gains from Simulations

In this section, we derive predictions of a simplified version of the imperfect foresight model of Gabaix and Laibson (2017) for our setting.

Hypothesis 3: Consider a decision-maker (DM) chooses between (u_0^E, u_1^L) and (u_0^L, u_1^E). For all $i \in \{E, L\}$, u_0^i, is received immediately and u_1^i is received in the following period. The DM knows the value of u_0^i with certainty, but lacks perfect information on the 'true' value of u_1^i and must generate simulations to forecast it. In the context of the first two studies, let (u_0^L, u_1^E) represent the utility from choosing to have only leisure time in WP1 such that all effort tasks are allocated to WP2, and (u_0^E, u_1^L) represent the utility of having only leisure in WP2 such that all effort tasks are allocated to WP1. We consider the case where the DM faces a tradeoff of allocating tasks to WP2, such that she has to do more total tasks when she chooses (u_0^L, u_1^E) than (u_0^E, u_1^L). In turn, u_1^L corresponds to only leisure in WP2 while u_1^E corresponds to unpleasant effort tasks in WP2, such that $u_1^L > u_1^E$. We refer to the former as the patient choice and the latter as the impatient choice.

Following Gabaix and Laibson (2017), normalize the DM's prior on u_1^i to zero such that $u \sim N(0, \sigma_u^2)$. This can be interpreted as the average utility that could be realized in WP2 given the choice set available to the DM. We consider the case where waiting periods prompt additional simulations relative to when no waiting periods. When the DM performs her first simulation of u_1^i, she draws an unbiased signal of its value $s_{1,1}^i = u_1^i + \epsilon_{1,1}$, where the first term in $(1, 1)$ corresponds to the time horizon and the second to the order of the signal drawn. The simulation noise $\epsilon_{1,1}$ is drawn from $\epsilon_1 \sim N(0, \sigma_{\epsilon_1}^2)$. Since we only consider a one-period time horizon, $\sigma_{\epsilon_1}^2 = \sigma_\epsilon^2$.[34] As a Bayesian, she integrates this signal with her prior. The DM's posterior forecast of u_1^i can be represented as $Ds_{1,1}^i$, where $D = \frac{1}{1+\sigma_\epsilon^2/\sigma_u^2}$. Integrating over the distribution of signals, we get $E_1(u_1^i) = Du_1^i$. Note that the expectation after one simulation is decreasing in the variance of the signal distribution σ_ϵ^2, which Gabaix and Laibson (2017) argue is a function of cognitive sophistication and information processing capacity.

[34] This is consistent with a one-period simulation under the proportional variance assumption of Gabaix and Laibson (2017), $\sigma_{\epsilon_t} = t \cdot \sigma_\epsilon$.

Before the initial simulation, the DM values the patient choice as u_0^E and the impatient choice as u_0^L. After the first simulation, the DM values the patient choice as $u_0^E + Du_1^L$ and the impatient choice as $u_0^L + Du_1^E$. Because $u_1^L > u_1^E$, it is straightforward to show that the DM's valuation of the patient choice increases after the initial simulation.

To illustrate how successive simulations increase the valuation of the patient choice relative to the impatient choice, let the DM draw a second signal $s_{1,2}^i = u_1^i + \epsilon_{1,2}$. She again updates her beliefs and obtains the posterior $Ds_{1,1}^i + Ds_{1,2}^i - D^2 s_{1,1}^i$. Integrating over the distribution of signals, we get

$$E_2(u_1^i) = Du_1^i + Du_1^i - D^2 u_1^i = 2Du_1^i - D^2 u_1^i \quad . \tag{4}$$

To illustrate the result, take a DM who is indifferent between the two choices after an initial simulation, $(u_0^E, u_1^L) \sim (u_0^L, u_1^E)$, such that her forecasted utility in expectation can be represented as

$$u_0^E + E_1(u_1^L) = u_0^L + E_1(u_1^E) \tag{5}$$

After the second simulation, the left hand side becomes $u_0^E + E_2(u_1^L)$ and the right hand side becomes $u_0^L + E_2(u_1^E)$. The change in valuation of the patient choice after the additional simulation can be calculated by subtracting the left side of (7) from (6) $+u_0^E$; the valuation change of the impatient choice can be calculated by subtracting the right side of (7) from (6) $+u_0^L$. The first term cancels out in both cases. The second term corresponds to:

$$E_2(u_1^i) - E_1(u_1^i) = 2Du_1^i - D^2 u_1^i - Du_1^i = Du_1^i - D^2 u_1^i = Du_1^i(1 - D) \tag{6}$$

The relative change in valuation of the patient choice to the impatient choice is captured by the difference $Du_1^L(1-D) - Du_1^E(1-D) = D(1-D)(u_1^L - u_1^E)$. Because $D \in (0, 1)$, the expression is positive.

Given this updating process, with $D \in (0, 1)$ and $(u_1^L - u_1^E) > 0$, it is straightforward to show that the successive simulations lead to a (weakly) greater valuation of the patient choice relative to the impatient choice. Under the assumption that a DM is prompted to generate more simulations with a waiting period than without one, this illustrates our third hypothesis.

Hypothesis 5: We now examine the relationship between the change in the forecasted value of

u_1^i after simulations and the variance of the simulation noise σ_ϵ^2.

As shown above, the relative change in valuation of the patient choice to the impatient choice after the initial simulation is $[(u_0^E + Du_1^L) - u_0^E] - [(u_0^L + Du_1^E) - u_0^L] = D(u_1^L - u_1^E)$. Because D is decreasing in the variance of the simulation noise, σ_ϵ^2, if the waiting period prompts the initial simulation relative to the case where the DM does not have a waiting period, then those with lower σ_ϵ^2 will exhibit greater increases in patience relative to those with higher σ_ϵ^2. Given the hypothesized relationship between σ_ϵ^2 and information processing capacity, this translates to waiting periods leading to greater increases in patience for those who score higher on this capacity measure — our Hypothesis 5.

However, because signals are unbiased and the true value of u_1^i does not differ by information processing capacity, successive simulations past the initial one lead to a convergence in patience across the different levels of information processing capacity. For example, consider the second simulation. Take the derivative of (8) with respect to D:

$$\frac{\partial(E_2(u_1^i) - E_1(u_1^i))}{\partial D} = u_1^i(1 - 2D) \quad . \tag{7}$$

This is negative if $D > \frac{1}{2}$.

$$D > \frac{1}{2} \Rightarrow \frac{1}{1 + \sigma_\epsilon^2/\sigma_u^2} > \frac{1}{2} \Rightarrow 1 + \frac{\sigma_\epsilon^2}{\sigma_u^2} < 2 \Rightarrow \frac{\sigma_\epsilon^2}{\sigma_u^2} < 1 \Rightarrow \sigma_\epsilon^2 < \sigma_u^2 \quad . \tag{8}$$

Because D is an inverse function of σ_ϵ^2, the expression (9) being negative implies that the change in forecasted value after the second simulation increasing in the variance of simulation noise.[35] Although DMs higher in information processing capacities start out more patient after the initial simulation, the patience of those who score lower on this measure converges to those who score higher with successive simulations. This implies that longer waiting periods lead to increasingly similar returns for patience across different levels of information processing capacity.

[35]Is $\sigma_\epsilon^2 < \sigma_u^2$? Gabaix and Laibson (2017) show that the *as-if* discount function produced by this model when the variance of forecasts of future utility is linearly increasing in the time horizon is hyperbolic with $\frac{\sigma_\epsilon^2}{\sigma_u^2}$ playing the role of the hyperbolic discounting parameter: $D(t-k) = \frac{1}{1+(t-k)\sigma_\epsilon^2/\sigma_u^2}$. Estimates of this parameter in the literature are considerably below one. In a large, representative sample, Andersen et al. (2014) estimate a hyperbolic discounting parameter of 0.089 (S.E. = 0.007), for example. Gabaix and Laibson (2017) parameterize their example with a value of 0.10.

A.2 Structural Estimation in the Online Effort Allocation Study

In this section, we discuss estimates of the utility parameters from equations (1), (2) and (3). Since participants make allocation decisions between two periods, each treatment on its own only reveals their one-hour discount factor for task effort. Because the timing of the work periods and the allocation decision differs by treatment, the variation in the theoretical interpretation of that discount factor allows us to identify the parameters of interest. Specifically, the treatments were designed to separately identify aggregate estimates of the exponential discount factor δ, the present bias parameter β, and the simulation parameter $S_k(t)$. The parameter $S_k(t)$ is meant to capture the effect of additional simulations of the decision problem prompted by the waiting period. In the application of the Gabaix and Laibson (2017) framework outlined in Section 2.1.2, the parameter can be represented as $S_k(t) = \frac{D_k(t)}{D_{k+1}(t)}$. Given the short horizon and the fact that our experiment manipulates the waiting period over only one interval, we drop the subscripts for the analysis, setting $S_k(t) = S$.

Our identification strategy is as follows. Participants in the Waiting Period treatment solve the optimization problem in equation (2), as laid out in Section 2.1.2. We allow for present bias, such that the discount factor between periods is equal to $\frac{D_1(1)}{D_1(0)} = \beta\delta$. Participants in the Immediate treatment solve a similar problem, shifted back by one period as in equation (1). The parameter S identifies any additional discounting that occurs in the Immediate treatment that does not occur in the Waiting Period treatment. Therefore, the discount factor in the Immediate treatment can be represented as $\frac{D_0(1)}{D_0(0)} = S\beta\delta$. We obtain an estimate of S as the ratio of the Immediate discount factor to the Waiting Period discount factor.

Participants in the Commit treatment maximize equation (3). At $t = 0$, subjects allocate tasks between $t = 1$ and $t = 2$. Because choices are made in the absence of a waiting period, the discount factor can be represented as $\frac{D_0(2)}{D_0(1)} \approx S\delta$.[36] In turn, we obtain an estimate of β as the ratio of the Immediate discount factor to the Commit discount factor.

Call z_1 tasks allocated to Work Period 1 and z_2 tasks allocated to Work Period 2 and r the the

[36]This is an approximation. Since the variance in forecasts of future utility is increasing in their time horizon, the as-if discounting that occurs in the Commit treatment is between one and two periods in the future, whereas in the Immediate treatment, it is between one period in the future and the present, which is subject to no uncertainty. Assuming a linear increase in simulation variance and time period, which leads to a hyperbolic as-if discount factor, the S in Commit is slightly closer to one than the S in Immediate. Our estimate of β is thus a lower bound on the quasi-hyperbolic discount factor.

interest rate by which undone tasks grow. The general convex intertemporal allocation decision in our study is

$$\min_{z_1,z_2} U(z_1,z_2) = z_1^{\gamma} + \delta_T z_2^{\gamma} \quad \text{s.t.} \quad z_1 + \frac{z_2}{1+r} = 40 \quad . \tag{9}$$

γ is the instantaneous disutility of effort parameter, and δ_T is a treatment-specific discount factor, which we map to the parameters of interest with the across-treatment comparisons mentioned above.

We make two additional adjustments to allow for more flexibility in our model of effort cost. First, we add background parameters ω_1 and ω_2 to the tasks required in each period to represent other effort that might need to be expended during those time periods. Second, we allow for the possibility of less-than complete recovery after Work Period 1 with another background effort parameter, ω_3, that enters as a coefficient on z_1 in the Work Period 2 effort level. The utility function is thus

$$U(z_1,z_2) = (z_1+\omega_1)^{\gamma} + \delta_T(z_2+\omega_2+\omega_3 z_1)^{\gamma} \quad . \tag{10}$$

We use the solution to the utility maximization problem to set up a maximum-likelihood estimation. The supply of tasks in Work Period 1 is

$$z_1^* = \frac{40A(1+r)+\omega_2 A-\omega_1}{1+A(1+r)-\omega_3 A} \quad , \tag{11}$$

where $A = \left(\delta_T(1+r-\omega_3)\right)^{\frac{1}{\gamma-1}}$. Individuals, i, solve this problem for each choice, j, and select the nearest available option subject to a standard normal error term, $\epsilon_{i,j}$, such that

$$z_{1,(i,j)} - \frac{40A_t(1+r_j)+\omega_2 A_j-\omega_1}{1+A_j(1+r_j)-\omega_3 A_j} + \epsilon_{i,j} = 0 \quad , \tag{12}$$

where $z_{1,(i,j)}$ is our observed choice for period 1 tasks by person i on task j. The likelihood associated with that observation is

$$\phi\left(z_{1,(i,j)} - \frac{40A_j(1+r_j)+\omega_2 A_j-\omega_1}{1+A_j(1+r_j)-\omega_3 A_j} \right) \tag{13}$$

When subjects select corner solutions from the convex choice sets, the convex first order conditions may poorly approximate choices. Therefore, we assume censoring at each corner as in a

Tobit model. If $z_{1,(i,j)} = 0$, then we assume that

$$\epsilon_{i,j} > \frac{40A_j(1+r_j) + \omega_2 A_j - \omega_1}{1 + A_j(1+r_j) - \omega_3 A_j} \quad , \tag{14}$$

and the likelihood contribution is

$$\Phi\left(- \frac{40A_j(1+r_j) + \omega_2 A_j - \omega_1}{1 + A_j(1+r_j) - \omega_3 A_j} \right) \quad . \tag{15}$$

If $z_{1,(i,j)} = 40$, then we assume that

$$\epsilon_{i,j} < \frac{40A_j(1+r_j) + \omega_2 A_j - \omega_1}{1 + A_j(1+r_j) - \omega_3 A_j} - 40 \quad , \tag{16}$$

and the likelihood contribution is

$$\Phi\left(\frac{40A_j(1+r_j) + \omega_2 A_j - \omega 1}{1 + A_j(1+r_j) - \omega_3 A_j} - 40 \right) \quad . \tag{17}$$

In our two binary choice tasks, subjects simply select the smaller value between $(40+\omega_1)^\gamma$ and $\delta_T(40(1+r)+\omega_2)^\gamma$. We make the standard Probit model assumption that the difference between the two utilities is subject to a normal distribution. Thus the probability of observing all work in the first period is

$$Pr(z_{1,(i,j)} = 40) = Pr((40+\omega_1)^\gamma - \delta_T(40(1+r_j)+\omega_2)^\gamma + \epsilon_{i,j} < 0) = \\ \Phi(\delta_T(40(1+r_j)+\omega_2)^\gamma - (40+\omega_1)^\gamma) \tag{18}$$

and the probability of observing all work in the second period is

$$Pr(z_1 = 0) = Pr((40+\omega_1)^\gamma - \delta_T(40(1+r_t)+\omega_2)^\gamma + \epsilon_{i,j} > 0) = \\ \Phi((40+\omega_1)^\gamma - \delta_T(40(1+r_j)+\omega_2)^\gamma). \tag{19}$$

These probabilities are used to construct the likelihood function. In the estimation, we impose the restrictions that $\gamma > 0$ and that $\omega_1, \omega_2, \omega_3 > 0$ to prevent degenerate results.

We estimate $\gamma = 1.255$ (S.E. = 0.047), indicating increasing marginal disutility of performing the counting task. There is no evidence on any background effort level in Work Period 1 ($\omega_1 = 0$), but there is evidence of background effort in Work Period 2 ($\omega_2 = 4.711$, S.E. = 2.455). Additionally there is some evidence of effort spillover across period ($\omega_3 = 0.253$, S.E. = 0.017). We estimate discount factors of $D_I = 0.968$ (S.E. = 0.071), $D_W = 1.151$ (S.E. = 0.113), $D_C = 1.062$ (S.E. = 0.087), and $D_{DC} = 0.878$ (S.E. = 0.061). They very short time horizon means that we should expect no discounting from an exponential model point of view. Only the Delay Control estimate is significantly different from one ($p = 0.65, 0.18, 0.48$, and 0.05 respectively). Estimates of S and β are discussed in the text.

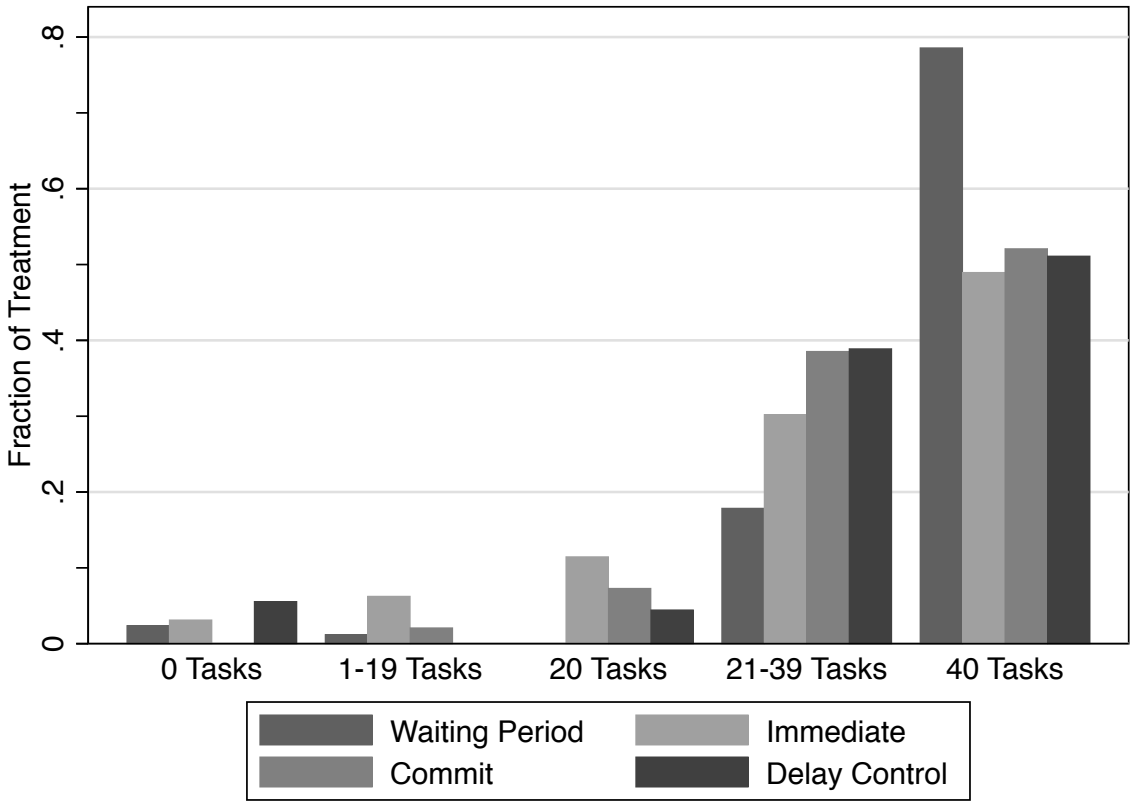

Figure A1: Task Allocations to Work Period 1 on Positive-interest Convex Choices, Online Study

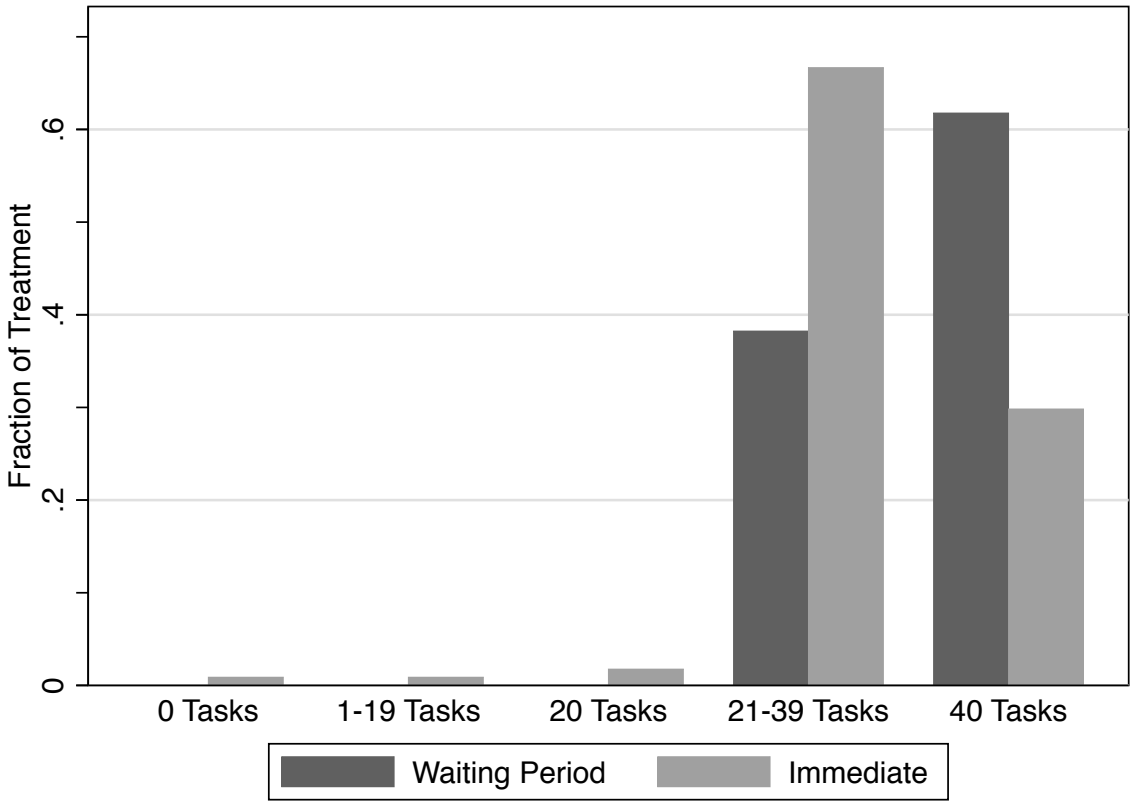

Figure A2: Task Allocations to Work Period 1 on Positive-interest Convex Choices, Lab Study

Table A1: Effect of Treatment on Convex Task Allocations to Work Period 1 Online Study, Tobit Model of Intensive and Extensive Margin Effects

Interest rate:	50%	25%	12.5%	0%
	(1)	(2)	(3)	(4)
Panel A: Treatment Effect Conditional on being Un-censored				
Waiting Period	3.967*	4.098*	3.822*	2.047
	(2.102)	(2.371)	(2.282)	(1.659)
Commit	0.333	0.852	1.124	-0.079
	(1.440)	(1.395)	(1.374)	(1.369)
Delay Control	0.152	0.120	0.181	-2.768*
	(1.541)	(1.533)	(1.498)	(1.458)
Constant	25.143	24.000	23.000	22.316
	(1.447)	(1.488)	(1.190)	(1.413)
Panel B: Treatment Effect on Probability of Doing All Tasks in WP1				
Waiting Period	0.309*	0.314*	0.297*	0.140
	(0.149)	(0.163)	(0.160)	(0.103)
Commit	0.026	0.065	0.087	-0.005
	(0.111)	(0.103)	(0.102)	(0.094)
Delay Control	0.012	0.009	0.014	-0.189*
	(0.119)	(0.117)	(0.116)	(0.112)
Constant	0.531	0.469	0.469	0.375
	(0.091)	(0.090)	(0.089)	(0.086)
N	122	122	122	122

* : $p < 0.10$. Each panel represents a different marginal effect of a Tobit model estimation, relative to the Immediate treatment. Bootstrapped standard errors from 1000 replications are reported in parentheses below each estimate. Output is reproducible with a seed of 1. We do not study extensive effects margin effect on the other boundary because the frequency of doing no tasks in WP1 is very low.

Table A2: Effect of Treatment on Convex Task Allocations to Work Period 1
Lab Study, Tobit Model of Intensive and Extensive Margin Effects

Interest rate:	50%	25%	12.5%	0%
	(1)	(2)	(3)	(4)
Panel A: Treatment Effect Conditional on being Un-censored				
Waiting Period	3.461**	4.526***	3.260**	1.608
	(1.336)	(1.338)	(1.364)	(1.535)
Constant	26.500	26.483	25.538	22.462
	(0.756)	(1.143)	(0.620)	(0.852)
Panel B: Treatment Effect on Probability of Doing All Tasks in WP1				
Waiting Period	0.251***	0.310***	0.228**	0.102
	(0.087)	(0.080)	(0.091)	(0.095)
Constant	0.342	0.237	0.316	0.316
	(0.075)	(0.067)	(0.074)	(0.075)
N	72	72	72	72

*** : $p < 0.01$, ** : $p < 0.05$. Each panel represents a different marginal effect of a Tobit model estimation, relative to the Immediate treatment. Bootstrapped standard errors from 1000 replications are reported in parentheses below each estimate. Output is reproducible with a seed of 1. We do not study extensive margin effect on the other boundary because the frequency of doing no tasks in WP1 is very low.

Table A3: Effect of Treatment on Convex Task Allocations to Work Period 1 By Belief Bias Level
Lab Study, Tobit Model of Intensive and Extensive Margin Effects

Sample:	Waiting Period		All	
	(1)	(2)	(3)	(4)
Panel A: Treatment Effect Conditional on being Un-censored				
Information Processing	1.698	2.170**	-0.420	-1.013
	(1.24)	(0.952)	(0.849)	(0.643)
$r = 0\%$		0.236**		0.343***
		(0.094)		(0.079)
Waiting Period			1.118	-0.749
			(2.176)	(1.363)
Waiting Period X Information Processing			2.544	3.982***
			(1.737)	(1.405)
Constant	24.578	10.217	24.752	14.132
	(1.066)	(1.457)	(0.756)	(1.921)
Panel B: Treatment Effect on Probability of Doing All Tasks in WP1				
Information Processing	0.145	0.169**	-0.029	-0.059
	(0.100)	(0.079)	(0.058)	(0.037)
$r = 0\%$		0.018***		0.020***
		(0.005)		(0.003)
Waiting Period			0.077	-0.044
			(0.151)	(0.079)
Waiting Period X Information Processing			0.175	0.232***
			(0.114)	(0.085)
Constant	0.613	-0.033	0.502	-0.079
	(0.112)	(0.148)	(0.076)	(0.090)
Clusters	34	34	72	72
N	102	102	216	216

*** : $p < 0.01$, ** : $p < 0.05$. Each panel represents a different marginal effect of a Tobit model estimation, relative to the Immediate treatment. Standard errors are clustered by individual. We do not study extensive margin effect on the other boundary because the frequency of doing no tasks in WP1 is very low.

Table A4: Observable Balance across Treatments, DRC Study

Variable	Immediate	Waiting Period	Difference
Female	0.41	0.42	-0.01
Age	30.90	30.59	0.31
Secondary education or beyond	0.79	0.77	0.02
Has children	0.69	0.75	-0.05
Employed	0.44	0.39	0.06
Distance from city center (1-3 scale)	1.57	1.61	-0.04
Feels safe at home (1-4 scale)	2.34	2.53	-0.20*
Access to food (1-4 scale)	2.39	2.39	0.00
Access to clean water (1-4 scale)	2.40	2.29	0.11
Access to medical care (1-4 scale)	2.05	2.13	-0.08
Access to shelter (1-4 scale)	2.36	2.40	-0.04
Access to phone network (1-4 scale)	2.66	2.40	0.26*
Life got better last year (1-5 scale)	3.04	3.14	-0.10
Expects life better next yr. (1-5 scale)	3.72	3.73	-0.08
Not afraid to take risks (1-4 scale)	3.03	3.12	-0.09
Feels in control of life (1-4 scale)	2.32	2.23	0.08
Worries about future (1-4 scale)	2.74	2.88	-0.14
Plans for next week (1-4 scale)	3.10	3.13	-0.04
Trusts others (1-4 scale)	2.38	2.55	-0.17
Close to community (1-4 scale)	2.94	3.05	-0.11
Property damage due to conflict	0.46	0.50	-0.04
Direct exposure to violence during war	0.38	0.30	0.08

* : $p < 0.10$.

Table A5: Impact of Waiting Period on Likelihood of Minimum-value Coupon Redemption, Probit Model

	(1)	(2)
Waiting Period	-0.211***	-0.243***
	(0.074)	(0.082)
Food Access		-0.020
		(0.045)
Food Access X Waiting Period		0.013
		(0.066)
Distance from Store		0.037
		(0.058)
Distance from Store X Waiting Period		0.038
		(0.087)
Trust in Others		-0.003
		(0.039)
Trust in Others X Waiting Period		0.022
		(0.066)
Risk Tolerance		-0.020
		(0.042)
Risk Tolerance X Waiting Period		0.133*
		(0.077)
Constant	0.250	0.251
	(0.037)	(0.038)
N	258	252

*** : $p < 0.001$, * : $p < 0.10$. Estimates are the marginal effect from Probit models, relative to the immediate treatment and the mean levels of the control variables. Robust standard errors are reported in parentheses below each estimate. Food access, trust in others and risk tolerance are measured on 1-4 scales. Distance from the store is measured on a 1-3 scale. All control variables are de-meaned. We lose six observations with the addition of control variables due to incomplete survey responses.

Sample Task

To continue, please complete the two example tasks below.

0	0	0	1	1	0	1	0	1	0	1	1	1	1	0
1	0	0	1	0	1	0	0	0	1	1	1	1	1	1
1	0	0	0	0	0	0	0	0	0	0	0	0	1	1
1	1	1	0	0	0	1	1	0	1	0	0	0	0	1
1	1	1	1	1	0	1	0	1	1	1	1	1	0	1
0	1	0	0	0	0	1	1	0	0	1	1	1	1	0
1	0	0	1	0	0	0	0	1	0	0	1	1	0	1
0	1	1	0	1	0	1	0	1	0	1	0	1	0	1
1	0	0	1	1	0	1	0	0	1	0	0	0	0	1
1	1	1	0	1	1	1	0	0	1	1	1	1	1	1

Example Table 1

How many zeros are in the table above?

0	0	1	0	1	1	0	1	1	1	0	0	1	1	0
0	1	0	0	0	0	1	1	1	0	0	0	0	0	0
1	1	0	0	1	1	0	1	1	0	0	1	1	0	0
0	0	0	1	0	1	1	1	1	0	0	0	0	1	1
1	0	1	1	0	0	1	0	1	1	0	1	0	1	1
0	1	1	1	1	0	1	0	0	0	1	1	0	0	0
0	0	0	1	1	0	0	1	1	1	1	0	1	0	0
0	0	0	0	1	0	0	0	1	0	0	1	0	0	1
0	1	1	0	1	1	1	0	1	1	1	1	1	0	0
1	1	0	1	1	1	0	0	0	0	1	0	1	0	0

Example Table 2

How many zeros are in the table above?

>>

Immediate Treatment

To finish the study and earn your payment, you will be given a choice about how many tasks to do and when to do them. The study is broken up into two work periods. Work Period 1 will begin immediately after you make a choice and will last for approximately 1 hour. Work Period 2 begins directly after that and also lasts for approximately 1 hour.

You will be given a choice of how many tasks to do in each work period.

If you choose not to work during a work period or if you finish your tasks early within a work period, you have free time for the rest of the work period. The program is timed in such a way that you cannot advance to the next work period until the full hour has elapsed. Once you finish the tasks you chose to complete in a work period, you can spend the rest of the time in the period however you want. For example, you can open another browser window and surf the internet, read a book, study, etc. Once the one hour ends, the next work period will begin. Here you will work on the number of tasks you chose to complete in that period.

Once Work Period 2 is over, you will have one additional hour of free time. After this, we will give you a brief survey to complete.

>>

Powered by Qualtrics

Waiting Period Treatment

To finish the study and earn your payment, you will be given choices about how many tasks to do and when to do them. We will first describe the choices to you. You will then have an hour to think about the choices and are free to do whatever you want to pass the time at your lab station. The one restriction is that you cannot communicate with others.

At the end of the hour you will be asked to make choices of how many tasks to do in each of the two work periods. Work Period 1 will begin after you make a choice and will last for approximately 1 hour. Work Period 2 begins directly after that and also lasts for approximately 1 hour.

You will be given a choice of how many tasks to do in each work period.

For example, you may be faced with a choice between:

Option A) 10 tasks in Work Period 1 and 0 tasks in Work Period 2
Option B) 5 tasks in Work Period 1 and 6 tasks in Work Period 2
Option C) 0 tasks in Work Period 1 and 12 tasks in Work Period 2

If you choose not to work during a work period or if you finish your tasks early within a work period, you have free time for the rest of the work period. The program is timed in such a way that you cannot advance to the next work period until the full hour has elapsed. Once you finish the tasks you chose to complete in a work period, you can spend the rest of the time in the period however you want. For example, you can open another browser window and surf the internet, read a book, study, etc. Once the one hour ends, the next work period will begin. Here you will work on the number of tasks you chose to complete in that period.

Once Work Period 2 is over, you will have a brief survey to complete.

Sample Convex Budget

You have a variety of options for how many tasks you need to do in Work Period 1 (starting immediately after your choices) and Work Period 2 (starting in approximately 1 hour).

Each column in the table below represents one option. For example, Option A involves doing 0 tasks in Work Period 1 *and* 60 tasks in Work Period 2. Option B involves doing 4 tasks in Work Period 1 *and* 54 tasks in Work Period 2, etc.

PERIOD	TASK OPTIONS										
	A	B	C	D	E	F	G	H	I	J	K
WORK PERIOD 1 - IMMEDIATELY	0	4	8	12	16	20	24	28	32	36	40
WORK PERIOD 2 - IN 1 HOUR	60	54	48	42	36	30	24	18	12	6	0

Please examine the table above and choose your preferred option by putting the letter of the option into the box below.

My preferred option is: