CESifo WORKING PAPERS

9472
2021

December 2021

Cognitive Uncertainty in Intertemporal Choice

Benjamin Enke, Thomas Graeber

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 9472*

Cognitive Uncertainty in Intertemporal Choice

Abstract

This paper studies the relevance of cognitive uncertainty – subjective uncertainty over one's utility-maximizing action – for understanding and predicting intertemporal choice. The main idea is that when people are cognitively noisy, such as when a decision is complex, they implicitly treat different time delays to some degree alike. By experimentally measuring and manipulating cognitive uncertainty, we document three economic implications of this idea. First, cognitive uncertainty explains various core empirical regularities, such as why people often appear very impatient, why per-period impatience is smaller over long than over short horizons, why discounting is often hyperbolic even when the present is not involved, and why choices frequently violate transitivity. Second, impatience is context-dependent: discounting is substantially more hyperbolic when the decision environment is more complex. Third, cognitive uncertainty matters for choice architecture: people who are nervous about making mistakes are twice as likely to follow expert advice to be more patient.

Keywords: cognitive uncertainty, intertemporal choice, complexity.

Benjamin Enke
Harvard University
Department of Economics
Cambridge / MA / USA
enke@fas.harvard.edu

Thomas Graeber
Harvard Business School
Boston / MA / USA
tgraeber@hbs.edu

December 3, 2021
Graeber thanks the Sloan Foundation for Post-Doctoral funding. Enke acknowledges funding from the Foundations of Human Behavior Initiative and Mind, Brain and Behavior at Harvard.

1 Introduction

Many economically important intertemporal decisions are complex and cognitively difficult. For instance, suppose your bank offers you a special promotional product, in which for every \$100 you invest, you receive \$105 back after three years. Which fraction of your income should you invest into this asset to maximize your discounted utility? Thinking through decisions like these is arguably challenging and may make people's decision process cognitively noisy.[1] This paper empirically examines how such noisiness impacts intertemporal choice. Our main hypothesis is that cognitive noise induces a "compression effect:" people behave as if they treat different time delays to some degree alike. Continuing the example, imagine an alternative scenario in which the promotional product pays back \$105 after five rather than three years. Our hypothesis is that cognitive noise induces people to take similar decisions across these two scenarios – a compression effect – even when their discounted-utility maximizing decision strongly varies.

This perspective has four main economic implications, for both theoretical modeling and applied work. First, extreme discounting over short horizons partly results from bounded rationality. If people implicitly treat a one-week delay as if it were longer, they will *appear* excessively impatient. Second, cognitive noise and the resulting compression logic shed light on various core empirical regularities that have proven difficult to explain (Ericson and Laibson, 2019; Cohen et al., 2020), such as why people appear much more impatient (per unit of time) over short horizons than over longer ones, why discounting is often hyperbolic even when the present is not involved, and why choices frequently violate transitivity. Third, a cognitive noise account predicts that phenomena that are of practical interest to economists, such as short-run impatience and hyperbolic discounting, should depend on decision complexity and opportunities for deliberation (Imas et al., 2021). Fourth, viewing intertemporal choice through a lens of cognitive noise has potential implications for choice architecture: to the degree that people are aware of their noisiness and resulting mistakes, they may welcome expert advice. In summary, we hypothesize that cognitive noise is crucial for understanding many economic aspects of intertemporal choice, in a way that is largely orthogonal to behavioral economists' traditional emphasis on taste-based present bias.

To empirically investigate these hypotheses, we use the insight that people often have some awareness of how noisy and prone to error their decision process is. *Cognitive uncertainty* refers to a decision-maker's subjective uncertainty over what their utility-maximizing decision is (Enke and Graeber, 2020). In an intertemporal choice

[1] Various lines of prior work posit the presence of cognitive noise. For example, in both drift-diffusion models (e.g, Krajbich et al., 2012; Fudenberg et al., 2018) and Bayesian noisy cognition models (Gabaix and Laibson, 2017; Woodford, 2020; Khaw et al., 2021; Frydman and Jin, 2021), agents accumulate noisy cognitive evidence.

context, such cognitive uncertainty could reflect (i) uncertainty over one's current or future preferences; or (ii) the cognitive difficulty of computing discounted utility conditional on knowing one's preferences, for example because the environment is complex, or because it is mentally hard to integrate one's preferences with payouts and delays.

In our experiments, we link discounting behavior to a quantitative measure of cognitive uncertainty. In a first paradigm, experimental participants make decisions in standard multiple price lists to trade off different time-dated UberEats vouchers that can be used for restaurant delivery and takeout. In a second, complementary paradigm, we implement analogous decisions, except that these are defined over hypothetical monetary amounts. Third, as a robustness check, we replicate our findings using a direct elicitation technique that does not have a response grid. We discuss in detail how our study design relates to discussions about experimental intertemporal choice methodology, including reliability and fungibility of payments.

Following Enke and Graeber (2020), after each choice, we elicit cognitive uncertainty as a person's subjective probability that their revealed valuation range of a larger-later payment (the switching interval in a choice list) actually contains their true valuation. We interpret this measure as capturing the participant's posterior uncertainty about their utility-maximizing decision, after a "cognitive signal" has been generated through deliberation. We document that measured cognitive uncertainty is significantly correlated with across-trial variability in responses to repetitions of the same choice problem. This suggests that cognitive uncertainty indeed captures a choice signature of cognitive noise, and that it is to some extent reflective of actual noisiness.

In our data, 75–80% of all decisions are associated with strictly positive cognitive uncertainty. The main insight of our analysis, from which many of our results follow, is that the choices of cognitively uncertain participants are strongly compressed in the sense that people value a payment delayed by one year very similarly to a payment delayed by two years. As a result of this compression, cognitively uncertain decisions *look like* they reflect high impatience over short horizons but low impatience over very long ones (a "flipping" property). Thus, cognitive uncertainty is strongly linked to the empirical regularity of decreasing impatience (hyperbolicity). Importantly, we find that cognitive uncertainty is predictive of short-run impatience and hyperbolic discounting also when the present is not involved in a tradeoff, such that models of present bias do not apply. All of these correlations are quantitatively large. For instance, the magnitude of decreasing impatience is five times larger for decisions that are associated with strictly positive cognitive uncertainty.

The idea that cognitively uncertain decisions look like different delays are treated alike also manifests in another widely-studied empirical regularity: subadditivity, a canonical transivity violation according to which people are more impatient over two short intervals than over one (normatively identical) long interval. As we predicted in our pre-registration, this pattern is likewise strongly correlated with cognitive uncertainty. Indeed, in many specifications, we cannot statistically reject the null hypothesis that there is no subadditivity in the set of cognitively certain decisions.

Our main hypothesis is that cognitive noisiness leads to an inelasticity of decisions with respect to the delay, rather than that it "replaces" taste-based present bias. Thus, as a placebo exercise, we pre-registered the prediction and find empirically that cognitive uncertainty is *unrelated* to front-end delay effects. These refer to the pattern that people tend to be more impatient about a delay that begins now rather than in the future, which is widely believed to be a revealed preferences signature of present bias. These results show that bounded rationality in the form of cognitive noise and taste-based present bias are separate, complementary objects that explain partly different phenomena.

We discuss in some detail how our results relate to the predictions of different classes of random choice models. Our results are consistent with models that represent average decisions as a convex combination of utility-maximization and a form of central tendency effect. For instance, in random response models (Regenwetter et al., 2018), this central tendency could reflect that the decision-maker plays randomly with some probability. In models of Bayesian noisy cognition (Gabaix and Laibson, 2017; Woodford, 2020), the central tendency could reflect an intermediate "cognitive default" action. For instance, because people have little experience with the decision context in our experiments, they may intuitively anchor on an "intermediate valuation" of the larger-later payment, as in numerous documentations of central tendency or compromise effects effects in psychology (Simonson and Tversky, 1992; Xiang et al., 2021). To illustrate this idea, we present a stylized Bayesian noisy cognition model of intertemporal choice in Section 3 and derive our pre-registered predictions within this framework.

To examine the quantitative importance of cognitive uncertainty, we estimate a model in which measured cognitive uncertainty determines the relative weights of utility maximization and the central tendency effect. In these estimations, accounting for cognitive uncertainty produces an increase in model fit that is twice as large as the increase resulting from incorporating present bias. At the same time, we also find that present bias is needed to rationalize the data, even when cognitive uncertainty is accounted for. This again suggests that present-biased preferences and cognitive noise are separate objects that both matter.

The reader may wonder why it is important for economists to understand that discounting behavior is to a large degree governed by bounded rationality rather than non-standard discount functions per se. After all, variants of the generalized hyperbola often fit data reasonably well, even if they may be getting the underlying reason wrong. We document that a cognitive noise and cognitive uncertainty account has real implications for economics, regarding both the intensive and extensive margins of decision-making.

First, on the *intensive margin*, an account of cognitive noise predicts systematic context-dependence of impatience: higher decision complexity, or higher cognitive busyness, should produce more pronounced hyperbolic discounting. According to stable discount functions, on the other hand, behavior does not vary as a function of complexity or cognitive states. To test these ideas, we implement additional treatment arms in which we either induce cognitive load or manipulate the complexity of the intertemporal decision tasks by embedding a math problem into them. As predicted, we find that both cognitive load and increased complexity lead to discounting behavior that is more hyperbolic. For example, subjects who were distracted by an addition task displayed more impatience over short horizons, but *less* impatience over very long ones.

Second, on the *extensive margin*, a main implication of our account is that people are potentially "nervous" about making impatient mistakes and might therefore desire expert advice. In contrast, in pure preferences-based accounts, people may behave in impatient ways, but at the time the decision is taken they do not worry that the decision reflects a mistake. We study this distinction by surprising participants with information about advice from a poll of economists, who recommend that the participant chooses the most patient available action. We find that the probability of revising the previously-taken decision in the direction of greater patience is twice as high among cognitively uncertain participants. We interpret these patterns as suggesting that cognitive uncertainty is relevant also for choice architecture.

Linking this paper to the literature, it is well-known that a broad class of anomalies appear to reflect an insensitivity of behavior with respect to the time delay. Various pioneering contributions to the literature focused on identifying (or even openly reverse-engineering) reduced-form discount functions that fit this compression pattern well (Mazur, 1987; Loewenstein and Prelec, 1992; Ebert and Prelec, 2007). Our approach builds on this work, but shows that these phenomena largely reflect bounded rationality. While we do not intend to claim that taste-based present bias is unimportant, we do argue that cognitive noise better explains many of the phenomena that are often ascribed to present bias, such as extreme short-run impatience and hyperbolicity of the discount function (but not, for example, dynamic inconsistency).

A second line of work in both economics and psychology has documented various "cognitive effects" on people's intertemporal decisions.[2] Most directly related to our work are models of cognitive or decision noise. A first class of models builds on the recent Bayesian cognitive noise literature (Woodford, 2020; Khaw et al., 2021; Gabaix, 2019;

[2]In economics, this includes work on time perception (Brocas et al., 2018), waiting periods (Imas et al., 2021), focusing effects (Dertwinkel-Kalt et al., 2021), similarity (Rubinstein, 2003), cognitive ability (Dohmen et al., 2010), and noise in commitment demand (Carrera et al., 2019).

Frydman and Jin, 2021; Frydman and Nunnari, 2021). Intertemporal choice applications of these models (Gabaix and Laibson, 2017; Gershman and Bhui, 2019) motivate our experiments, though as we discuss below they do not fully explain our evidence. A second class of models comprises random response and random preference models, which have received attention in a largely theoretical literature in decision theory and mathematical psychology (e.g., Regenwetter et al., 2018; Lu and Saito, 2018; He et al., 2019).[3] Relative to these various strands of the literature, our contribution is that we directly measure and exogenously manipulate cognitive noise, which allows us to provide much sharper and more direct tests, and to quantify the relative contribution of cognitive noisiness to intertemporal choice.

The idea of empirically measuring cognitive noise and related concepts is increasingly gaining traction in the economics literature (Butler and Loomes, 2007; Agranov and Ortoleva, 2017, 2020; Khaw et al., 2021; Enke and Graeber, 2020). While economists have not deployed these techniques in intertemporal choice contexts, two psychological and neuroscientific papers that are contemporaneous with ours also elicit people's confidence in their intertemporal decisions (Bulley et al., 2021; Soutschek et al., 2021). While their objectives, approaches and measurement methods differ from ours in many ways, probably the biggest difference is that they do not focus on the objects that we are interested in here: explaining widely-studied empirical regularities from the economics literature, documenting a dependence of choices on complexity, structurally estimating a cognitive noise model, and highlighting implications for choice architecture.

Finally, abstracting away from intertemporal choice as such, our work ties into an active recent literature that suggests that what seem like non-standard preferences are sometimes better thought of as reflecting bounded rationality (Esponda and Vespa, 2016; Nielsen and Rehbeck, 2020; Imas et al., 2021; Martínez-Marquina et al., 2019; Bordalo et al., 2020) and complexity (Abeler and Jäger, 2015; Oprea, 2020).[4]

The paper proceeds as follows. Section 2 provides a brief review of the extant evidence on compression effects in intertemporal choice. Section 3 discusses theoretical background and develops our predictions. Section 4 presents the experimental design, Sections 5–6 the results and Section 7 the model estimations. Section 8 shows results on choice architecture and exogenous manipulations of complexity. Section 9 concludes.

[3]As summarized by Regenwetter et al. (2018), researchers in this (mostly psychology) literature typically argue their empirical case through model-fitting exercises rather than direct measurements of noisiness. A limitation of this approach is that with a multitude of different models as well as functional form specifications at the researcher's disposal, a large set of different results can potentially be rationalized.

[4]Our focus on *cognitive* uncertainty also links to the "implicit risk" literature, which highlights the importance of *objective* uncertainty about whether or when a delayed reward is received (Sozou, 1998; Dasgupta and Maskin, 2005; Halevy, 2008; Chakraborty et al., 2020).

2 Intertemporal Choice Regularities

The behavioral economics research program has successfully shed light on how taste-based present bias contributes to the stylized fact that people often appear extremely impatient and dynamically inconsistent. Yet, as highlighted by recent review papers, other commonly-identified – and economically no less important – empirical regularities are less well-understood and not easily explained by accounts of present bias (Ericson and Laibson, 2019; Cohen et al., 2020). A key stylized fact is that people's discounting behavior tends to be very inelastic with respect to the length of the time delay. This basic principle manifests in three distinct empirical regularities. First, as visualized in Panel A of Figure 1, people often act in very impatient ways in decisions over relatively short horizons, yet appear considerably less impatient over longer horizons, in both lab and field (Thaler, 1981; Loewenstein and Prelec, 1992; Giglio et al., 2015). This implies that people's implied per-period impatience strongly decreases in the length of the time delay. This is puzzling because the extreme flattening out of observed discounting behavior is not predicted by present bias models (Laibson, 1997), which converge to exponential discounting over long horizons. Second, as visualized in Panel B of Figure 1, an inelasticity of discounting with respect to the time delay is also observed for tradeoffs in which the early consumption opportunity is not today but in the future, again at odds with a pure present bias account (Kable and Glimcher, 2010).[5] Third, experimental studies robustly identify a particular type of transitivity violation called subadditivity, according to which people appear considerably more patient in tradeoffs over one long interval than in choices where that same interval is partitioned into two sub-intervals (Read, 2001). Again, this can be understood as people being insensitive to the time delay. A dominant approach in the economics literature has been to attempt to explain these stylized facts through non-standard discount functions, such as the generalized hyperbola and its variants (e.g., Mazur, 1987; Loewenstein and Prelec, 1992; Kable and Glimcher, 2010). Either implicitly or explicitly, such accounts usually take the perspective that "anomalous" discounting behavior reflects "anomalous" preferences. We, on the other hand, study the hypothesis that they arise from cognitive noise.

3 Theoretical Considerations and Hypotheses

Consider a choice context in which a decision-maker (DM) is prompted to specify the units of consumption a at an earlier point in time t_1 that make him indifferent to con-

[5] Decreasing impatience is the dominant finding in the literature (see, e.g. Cohen et al., 2020; Kable and Glimcher, 2010; He et al., 2019). However, it is not universal, neither when the early date is today nor when it is in the future (see, e.g., Andersen et al., 2014; Harrison et al., 2005).

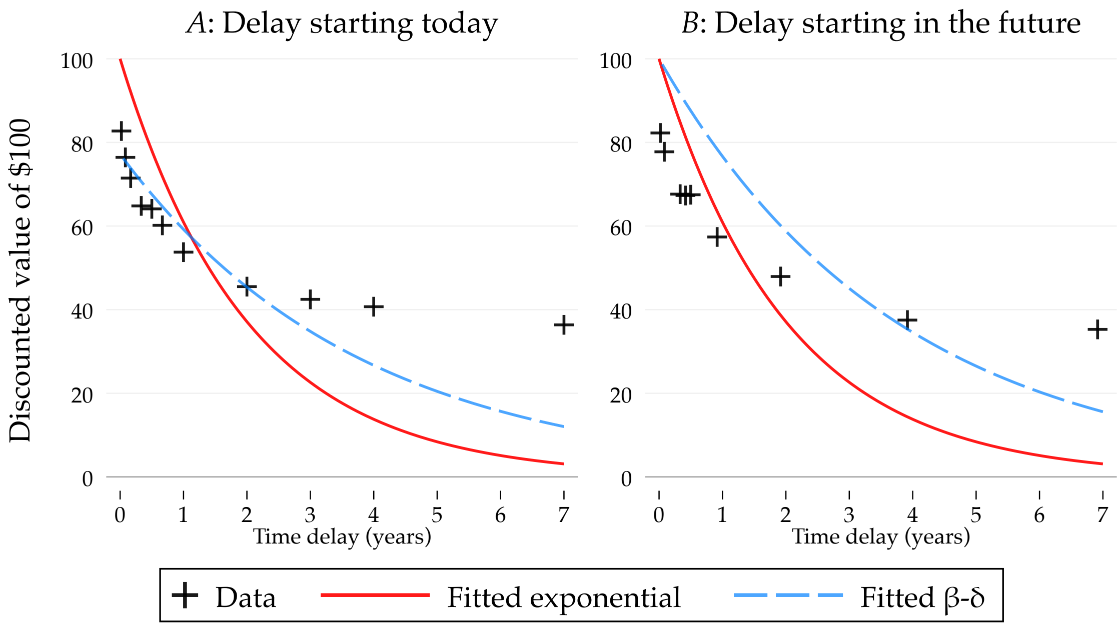

Figure 1: The figure shows the discounted value of $100 to be received with different time delays, partitioned by whether the early payment date is today (Panel A) or in the future (Panel B). The black markers indicate average behavior in our experiments described in Section 4. The red solid line fits an exponential discounting model and the blue dashed line a $\beta - \delta$ model. Both models are estimated on the joint data.

suming $c_{t_2} = 1$ at $t_2 > t_1$. We define $\Delta t \equiv t_2 - t_1$. Denote by $D(t) = \delta^t$ the DM's discount function, and by $u(\cdot)$ a weakly concave utility function[6]. Later, we will allow for taste-based present bias. A helpful theoretical benchmark is that of a rational DM's utility-maximizing decision, which equates the discounted utilities of both options. Normalizing $u(1) = 1$, we get:

$$D(t_1)u(a) = D(t_2)u(1) \quad \Rightarrow \quad a^* = u^{-1}(\delta^{\Delta t}) \in [0, 1]. \tag{1}$$

Our objective in this section is to summarize the extant random choice literature in a way that highlights that three broad classes of models often generate an inelasticity of decisions with respect to the delay. Per the discussion in the preceding section, this inelasticity is a potential driver behind various empirical regularities. We only provide a brief discussion here because, as reviewed in the formal treatment of Regenwetter et al. (2018), random choice models exhibit large diversity in precise modeling approaches and functional form assumptions.

Bayesian cognitive noise models. Bayesian cognitive noise (also called cognitive imprecision) models presume that people do not have direct access to their utility-maximizing action a^* but only to a noisy cognitive signal. In the intertemporal choice domain, cognitive noise could arise for a variety of reasons, which we discuss in Section 4 below. In

[6]When $u(c) = c^\alpha$, eq (1) also applies in the case $c_{t_2} \geq 1$ where a^* is now interpreted as the normalized indifference point of the rational DM.

general, cognitive signals can be conceptualized as pertaining to a problem parameter (Khaw et al., 2021; Frydman and Jin, 2021), utility from future consumption (Gabaix and Laibson, 2017), or the utility-maximizing action itself. We here adopt the latter approach. DMs are hypothesized to combine this cognitive signal with a prior over their utility-maximizing action, which we call cognitive default action (see, e.g., Woodford, 2020). This cognitive default can be interpreted as the action the DM would take in the absence of any deliberation.

Appendix A presents an intertemporal cognitive noise model that is an adaptation of atemporal applications (Fennell and Baddeley, 2012; Heng et al., 2020). In this model, the DM holds a Beta-distributed prior over his discounted-utility maximizing action, where the prior has mean d. Through deliberation, the DM generates a cognitive signal about what his discounted-utility maximizing decision is. This signal S is (scaled) Binomially distributed and satisfies $E[S] = a^*$. Given signal realization $S = s$, a Bayesian DM's posterior mean over his utility-maximizing action can be represented as

$$a^o = \lambda s(a^*(\delta, \Delta t)) + (1-\lambda)d \qquad \Rightarrow \qquad E[a^o] = \lambda a^*(\delta, \Delta t) + (1-\lambda)d. \quad (2)$$

This formulation intuitively captures an anchoring-and-adjustment heuristic (Tversky and Kahneman, 1974), according to which people anchor on some initial reaction (which may vary across contexts) and then adjust based upon the outcome of their deliberation process. Here, the weight λ partly captures the precision of the cognitive signal.[7]

The main implication of eq. (2) is that decisions are insufficiently sensitive to the time delay because they partly reflect the delay-invariant cognitive default. As we show in Appendix A, eq. (2) implies the following predictions, which we pre-registered.

Pre-registered predictions. *A DM with $\lambda < 1$ (cognitive noise) exhibits:*

1. *More pronounced short-run impatience, both when the time delay starts in the present and when it starts in the future.*

2. *More pronounced decreasing impatience, both when the time delay starts in the present and when it starts in the future.*

3. *More pronounced subadditivity.*

4. *The same degree of front-end delay effects: the pattern that people appear more patient when a constant is added to both the early and the later date.*

[7] Another potential microfoundation for regression to a cognitive default d are models of caution (e.g., Cerreia-Vioglio et al., 2015; Chakraborty, 2020). These models are intuitively related in the sense that – due to subjective uncertainty over their utility function – agents regress towards preferring a certain option. This setup could potentially be modified such that, because of uncertainty about their preferences and caution, agents regress to a "simple" or "intuitive" option d, rather than a certain one.

To see the logic behind these predictions, it is useful to imagine that the cognitive default action d in eq. (2) is somewhat "intermediate" in nature.[8] In our experiments, a central tendency is plausible because people have little experience with the context, which plausibly leads them to intuitively anchor on an "intermediate valuation" (away from $a^o = 0$ and $a^o = 1$), akin to documentations of central tendency or compromise effects effects in psychology. Then, Predictions 1 and 2 imply a distinctive "flipping" property: while cognitively noisy agents appear *more* impatient over short delays, the inelasticity with respect to the delay can make them *less* impatient over very long delays.

Prediction 4 clarifies that a cognitive noise framework like the one sketched above does not predict a link between cognitive noise and front-end delay effects, which are usually thought of as a canonical signature of present-biased preferences.[9] The reason is that, in eq. (2), cognitive precision λ only affects how people respond to a given time delay, rather than whether it starts in the present or future.

Finally, we note that not all cognitive noise models generate the full set of predictions above. The main intertemporal choice applications of Bayesian noisy cognition in the literature are Gabaix and Laibson (2017) and Gershman and Bhui (2019). Their setup is slightly different from the one above because they assume that all decision-relevant cognitive noise stems from the mental simulation of future *utils*. The stylized framework above, on the other hand, takes a broader perspective by assuming that the *utility-maximizing action* is simulated with noise, regardless of what the underlying sources of noise are (noisy mental simulation of future utils may be one of them). This distinction matters for predictions. While Gabaix and Laibson's model generates decreasing impatience, their model makes two predictions that differ from the ones above. First, because their model maintains transitivity, it does not predict subadditivity. Second, their model predicts front-end delay effects and related preference reversals, see Section 2.7 of Gabaix and Laibson (2017).

Random response models. This class of models is broad. One incarnation that relates to the preceding discussion is that the DM probabilistically either plays his utility-maximizing action or chooses at random, $\epsilon \sim F(\cdot) \in [0,1]$, with $E[\epsilon] = d$. Formally, we say that a trembling action a^{tr} is given by

$$a^{tr} = \begin{cases} a^*(\delta, \Delta t) & \text{with prob. } \lambda \\ \epsilon & \text{otherwise} \end{cases} \qquad \Rightarrow \qquad E[a^{tr}] = \lambda a^*(\delta, \Delta t) + (1-\lambda)d. \quad (3)$$

[8]Specifically, suppose that $a^*(\Delta t \to 0) > d > a^*(\Delta t \to \infty)$, which says that the default action is less (more) patient than the utility-maximizing action for very short (long) time delays.

[9]Front-end delay effects refer to the regularity that people generally behave less patiently in a tradeoff between consumption dates t_0 and t_1 than in a tradeoff between $t_0 + z$ and $t_1 + z$, for $z > 0$.

This expression for the average action is identical to the one in (2). Thus, the two models make identical predictions about average behavior. Moreover, the models are also difficult to tease apart looking at individual decisions because they both predict that actions will be random (even conditional on potential anchoring on a cognitive default). Thus, depending on the precise assumptions about the distribution of the random response, various different individual-level response patterns can be rationalized.[10]

Random preference models. This class of models assumes that the DM's discount function is stochastic and fluctuates over time (e.g., Regenwetter et al., 2018; Lu and Saito, 2018; He et al., 2019). In its most widespread incarnation, random intertemporal preferences models assume that "true" discounting is exponential, yet the decision-relevant discount factor $\tilde{\delta}$ varies randomly across trials, such that $\tilde{\delta} = \delta + \mu$, with $E[\mu] = 0$. Thus, in the setup sketched above, the DM's random preference action a^r would be given by

$$a^r = a^*(\tilde{\delta}, \Delta t). \tag{4}$$

It is widely understood that variation in δ can produce behavior that implies "decreasing impatience" because the average of multiple exponential functions is not necessarily exponential and can be hyperbolic. This insight was first noted in interpersonal (social welfare) contexts (Weitzman, 2001; Jackson and Yariv, 2014). However, a mathematically identical insight applies when a single DM's discount factor varies across time (Lu and Saito, 2018; He et al., 2019). Thus, as in the models described above, higher noisiness (variance of μ) should be correlated with stronger decreasing impatience. At the same time, in contrast to the cognitive noise model sketched above, models that only feature random variation in preferences *do* predict front-end delay effects (see Proposition 1 in Jackson and Yariv (2014)) and *don't* predict subadditivity.[11]

Summary and empirical implementation. The different classes of random choice models can make similar predictions, in particular as far as a link between noise and decreasing impatience are concerned. Moreover, the models afford varying degrees of flexibility (see Regenwetter et al., 2018). Hence, our objective is not to definitely tease

[10]Another type of random response model is that the DM's action is given by $a^{tr,2} = a^*(\delta, \Delta t) + \eta$, with $E[\eta] = 0$. Then, because the DM's action is bounded by zero and one, random decision errors may lead to "bouncing off the boundary" and push decisions to be intermediate, on average. We do not highlight this type of model because, in our data, decisions that are associated with strictly positive cognitive uncertainty are rarely located at or close to the boundaries, see Appendix Figure 13.

[11]To see this, consider a model à la Lu and Saito (2018), in which the DM draws a separate discount factor $\tilde{\delta} = \delta + \mu$ for each potential calendar time prior to observing the specific payout dates in an experimental trial. In a subadditivity documentation, there are three time delays, $t_0 \to t_1$, $t_1 \to t_2$ and $t_0 \to t_2$. Let the decision-relevant discount factors for the first two delays be $\tilde{\delta}_{t_0 \to t_1} = \delta + \mu_1$ and $\tilde{\delta}_{t_1 \to t_2} = \delta + \mu_2$. Then, with $\tilde{\delta}_{t_0 \to t_2} = \tilde{\delta}_{t_0 \to t_1} \tilde{\delta}_{t_1 \to t_2}$, there is no subadditivity.

these models apart, but to generically show that cognitive noise is instrumental for understanding intertemporal choice. At the same time, to the degree that the different classes of models *do* make different predictions, our empirical results will allow us to draw some conclusions about the relative explanatory power of the different approaches.

Because the actual form and realizations of cognitive noise are unobservable, we empirically measure a signature of cognitive noise. Following Enke and Graeber (2020), we use the language of cognitive noise models to define *cognitive uncertainty* as people's lack of certainty that their action equals their true utility-maximizing action:

$$p_{CU} \equiv P(\left|A|\{S=s\} - a^o\right| > c), \quad (5)$$

where $A\big|\{S = s\}$ is the perceived posterior distribution about the utility-maximizing action, conditional on the cognitive signal s. Intuitively, cognitive uncertainty captures the likelihood with which the DM thinks his optimal action might fall outside a window of length $2c$ around the action that he actually chose.

4 Experimental Design

4.1 Choice Tasks

Incentivized UberEats Voucher Experiments. In treatment *Voucher Main*, rewards are given by UberEats food delivery vouchers.[12] Participants complete multiple price lists (MPLs) that elicit interval information about indifference points. In each list, the left-hand side Option A is a fixed delayed UberEats voucher with value $y_2 \in \{40, 42, \ldots, 50\}$. The later payout date $t = t_2$ varies between one week and one year. The right-hand side Option B is an UberEats voucher the value of which increases as one goes down the list, from \$2 to \$$y_2$, in steps of \$2 each. The payment date for Option B, t_1, is always strictly earlier than the one for Option A, though not necessarily today.

Participants had to indicate a choice between Options A and B in each row of the MPL. We implemented a computerized auto-completion mode that enforces a single switching row: whenever a subject chose Option A in a given row, Option A automatically got selected in all rows above. Likewise, whenever a subject chose Option B in a given row, Option B automatically got selected in all rows below. Participants could revisit and change their choices at any time, and choices only became locked in when a participant decided to proceed to the next screen. Appendix Figure 10 shows a screenshot.

[12] The currently most widely used experimental economics paradigm to implement primary rewards in an intertemporal choice context consists of real effort tasks. These are infeasible in our context, however, because our research hypothesis requires a consumption good that can plausibly be implemented with long time delays, while real effort studies focus on horizons of a few weeks at most.

UberEats is the largest online food ordering and delivery service in the world. The service can be used to order food for takeout or delivery from a wide array of restaurants and is widely available throughout the United States, with an estimated market share of between one fifth to one third (Curry, 2021). Through a special collaboration with Uber, we designed our UberEats vouchers to be valid for a period of only seven days. For example, when a choice option is given by "$40 voucher that is valid in 6 months," then this means that the voucher will become valid six months after the participant's study date, and will remain valid for a period of seven days. We implemented a comprehension check to verify that participants understood that the voucher would expire after seven days, rather than be valid indefinitely. Participants' vouchers were directly credited to their personal UberEats accounts within 10 hours of completion of the study, such that subjects did not have to actively claim the voucher. The vouchers were always visible in their accounts, they could just not be used before the validity period. Participants received automatic reminders 24 hours before a voucher became valid and 24 hours before it expired.

Hypothetical Money-Early-versus-Later Experiments. Treatment *Money Main* has the same structure as the UberEats voucher experiments, except that the rewards are given by hypothetical dollar amounts. While the hypothetical nature of the payouts has obvious disadvantages, it also confers various advantages, in particular in conjunction with our financially incentivized UberEats experiments. First, we could explicitly instruct participants to make their choices assuming that there is no payment risk. Second, hypothetical payments allow us to use some very long time delays (up to "in 7 years") that would not be credible with real payments or food vouchers. This is an important advantage because, as discussed above, the inelasticity of discounting to the time delay leads us to expect that the relationship between cognitive uncertainty and impatience will flip as a function of the time delay. Finally, money experiments allow us to replicate the setup in which regularities such as diminishing impatience or subadditivity have predominantly been documented in the literature (Cohen et al., 2020).

Choice configurations. First, for choice lists with an early date of today, we implement delayed dates that range from one week to seven years in the hypothetical money experiments, and from one week to one year in the incentivized UberEats study. Second, in both experiments, we implement a broad set of lists that have an early payment date of "in one month," again with large variation in the corresponding later payment dates. These choice lists allow us to study short-run impatience and decreasing impatience, starting from both the present and the future.

Third, we implement sets of three choices each that serve to test for subadditivity effects, such as: $(t_1 = 0, t_2 = 12m)$, $(t_1 = 0, t_2 = 6m)$, $(t_1 = 6m, t_2 = 12m)$. Fourth, these subadditivity sets also allow for an analysis of front-end delay effects: the extent to which people are more patient in, e.g., $(t_1 = 6, t_2 = 12)$ than in $(t_1 = 0, t_2 = 6)$. Fifth, for each participant, two randomly selected choice configurations were presented twice in random locations in the sequence of twelve price lists. These are exact repetitions of the same choice problems and facilitate an analysis of across-trial choice variability. The order of all choice lists was randomized at the participant level.

Study components. The hypothetical money study consisted of four parts. In the first, each participant completed a total of twelve MPLs. In the second part, each subject completed six additional intertemporal choice problems that were administered in a direct elicitation format rather than using MPLs. We discuss these data in greater detail in Section 6.4. In the third part of the study, participants completed three choice under risk MPLs that (i) facilitate an analysis of the cross-domain stability of cognitive uncertainty and (ii) allow to disentangle time discounting from the role of utility curvature in our structural analyses (Section 7). In the fourth part, participants completed a Raven matrices test of cognitive skills. The structure of the UberEats study was identical, except that we did not implement the direct elicitation choice problems.

4.2 Measuring Cognitive Uncertainty

Elicitation. In both paradigms described above, participants make choices in MPLs that carry interval information about indifference points. In our experiments, the switching intervals have a width of $2. Our experimental instructions explain that we use this switching interval to determine how much the participant values the larger-later payment at the earlier date. Immediately after each choice list, we measure cognitive uncertainty (CU) as the participant's subjective probability that their true valuation of the later payment / voucher is actually contained in their stated switching interval. Specifically, after a participant completes a choice list with switching interval given by $[\$a, \$b]$, the subsequent screen reminds them of their previous decision and elicits cognitive uncertainty:

> *Your choices on the previous screen indicate that you value $\$y_2$ in t_2 somewhere between \$a and \$b in t_1. How certain are you that you actually value $\$y_2$ in t_2 somewhere between \$a and \$b in t_1?*

Participants answer this question by selecting a radio button between 0% and 100%, in steps of 5%. Appendix Figure 11 provides a screenshot. This cognitive uncertainty measurement follows the same protocol as proposed in a revised version of Enke and Graeber

(2020) for choice under risk, here adapted to an intertemporal choice context. In line with the discussion in Section 3, we interpret this question as capturing the participant's (posterior) uncertainty about their utility-maximizing decision, after some sampling of cognitive signals has taken place. We refer to (inverted) responses to this question as *cognitive uncertainty* rather than confidence because in economics the latter is used for problems that have an objectively correct solution.

Potential origins of cognitive uncertainty. Our measure is deliberately designed to capture participants' overall subjective uncertainty about what their preferred action is. This uncertainty could have various potential origins. First, people may not know their true preferences. This preference uncertainty could either be about one's true discount factor, or about the instantaneous utils that one will derive from future consumption, as in Gabaix and Laibson (2017).

Second, even conditional on knowing their preferences, people may cognitively struggle with choosing an action that maximizes discounted utility. For example, people may find it hard to cognitively integrate their discount factor with the time delay that is implied by different choice options, or they may suffer from imperfect time perception. A hypothetical special case of this class of non-preference-uncertainty mechanisms is that there is no true discounting at all, but that experimental subjects find it cognitively difficult to maximize the net present value of payments.

Comparison with alternative measures. Broadly speaking, the literature has proposed two different types of measures for eliciting people's uncertainty about their own decisions. At one extreme, psychologists, neuroscientists and some economists elicit measures of "decision confidence," in which subjects indicate on Likert scales how "confident" or "certain" they are in their decision (e.g., De Martino et al., 2013, 2017; Polania et al., 2019; Bulley et al., 2021; Xiang et al., 2021; Butler and Loomes, 2007). At the other extreme, economists have proposed to use measures of across-trial variability (Khaw et al., 2021) or deliberate randomization (Agranov and Ortoleva, 2017, 2020). Our preferred measure strikes a middle ground between these two approaches. While our approach retains the attractive simplicity of implementing a single question (as in the psychology literature), it is also quantitative in nature. The simplicity of asking one question per decision screen should be contrasted with the approach of gauging cognitive noise through across-task variability in choices, which requires *many* trials and is usually defined at the level of a study rather than of a single choice problem. Below, we document a strong correlation between our simple-but-unincentivized CU measure and choice variability.

4.3 Design Considerations

Time discounting studies are complicated by a range of methodological considerations. We discuss prominent concerns and implications for interpretation below.

External uncertainty / payment credibility. According to the so-called "implicit risk" hypothesis, intertemporal decisions could reflect not only genuine discounting but also external uncertainty (e.g. Benzion et al., 1989; Sozou, 1998; Halevy, 2008). This could be due to a lack of trust in the experimenter, uncertainty about the future purchasing power of money or vouchers, or the subjective probability of forgetting about the existence of the later reward. To address this, we put various measures in place. First, we deliberately implemented the money experiments in hypothetical terms. This allows us to emphasize that subjects should make their decisions assuming that they know with certainty that they will receive all payments as indicated. We verified understanding of this through a comprehension check.

Second, in the UberEats experiments, because vouchers appear in the participant's UberEats account within a few hours of the study regardless of the precise validity period, there is no differential payment risk across vouchers with different time delays. Participants could always view vouchers in their account, they could just not be used. We view this as a main advantage of our method relative to traditional monetary payments.

Third, those participants that actually won a voucher were asked to state their subjective probability that they will actually receive and use their voucher. The median (average) response is 95% (84%). Most importantly, we find that subjects' beliefs are uncorrelated with the delay of the voucher's validity period. This suggests that future vouchers were not perceived as more uncertain. All of our results are robust to only including participants in the analysis who indicate 100% certainty.[13]

Cognitive vs. external uncertainty. A related concern is that participants misinterpret the CU question as asking about their subjective probability of actually receiving the later reward. To address this, our money experiments include a comprehension check question that directly asks participants to indicate whether the CU elicitation question asks about (i) the subject's subjective probability of actually receiving the money or (ii) their certainty about own their valuation, given that they know they will receive the money with certainty. In addition, notice that an account of CU capturing perceived payment uncertainty would predict that CU is always negatively correlated with observed patience.

[13]Regarding actual consumption of our vouchers, at the time of the writing of this paper, 77% of subjects had used their UberEats credit, which is arguably a high usage rate for a voucher. This percentage fluctuates across delays but does not systematically decrease in the length of the delay.

However, we will see that, over sufficently long time horizons, CU is actually positively correlated with patience.

Fungibility. A common argument is that intertemporal choice experiments over money do not capture preferences-based discounting because money is fungible. From such a perspective, behavior in experiments reveals participants' attempt to maximize the net present value of payments, given (perceived) real interest rates. An alternative view is that experimental participants narrowly bracket their choices and treat monetary amounts in experiments as proxy for utils (Halevy, 2014; Sprenger, 2015; Andreoni et al., 2018; Epper et al., 2020). We acknowledge this discussion, but note that it only affects the precise interpretation of our cognitive uncertainty question. Under the interpretation that our experimental paradigms do not capture true discounting, our CU measure picks up participants' cognitive limitations in computing discounted utility (here: NPVs), conditional on knowing their preferences ($\delta = 1$). On the other hand, if experiments over money also capture real discounting, the CU question potentially captures all of the various psychological mechanisms discussed in the previous subsection. Regardless of whether the participant's objective is to maximize NPV or discounted utility more generally, our hypothesis is that subjective uncertainty about the utility-maximizing action is associated with a compression effect.

Utility curvature. Estimates of discount rates from price list choices may be confounded unless the curvature of the utility function is taken into account. To address this, we use the "double price list method" that estimates utility curvature from a separate set of risky choices.

Transaction costs. A main concern with traditional time discounting experiments is that they capture differential transaction costs between present and future. In our hypothetical money experiments, transaction costs are implausible. In the UberEats experiments, there are likewise no transaction costs because participants automatically receive their vouchers credited to their UberEats app, together with automated reminders about the validity dates.

4.4 Logistics and Participant Pool

The study was conducted on Prolific, an online worker platform. Recent experimental economics work suggests that data quality on Prolific is higher than on Amazon Mechanical Turk, and comparable to that in a canonical lab subject pool (Gupta et al., 2021). For the hypothetical money experiments, we made use of Prolific's "representative sample" option to collect data from a broad and diverse (though not actually nationally representative) set of participants.[14] We pre-registered a sample size of $N = 600$ participants. However, because of the discreteness of Prolific's representative sample procedure, we eventually ended up sampling $N = 645$ people. Since we view throwing away data as questionable, we keep the full sample, but we have verified that all results hold if we restrict the sample to the first 600 completes.

In the UberEats experiments, the study description that was visible to prospective participants announced that study bonuses would be paid in the form of UberEats vouchers. In addition, we implemented a screening in which participants were again asked whether they possess an UberEats account, and we immediately routed all people out of the experiment who did not.[15] As we pre-registered, $N = 500$ workers participated in the UberEats study.

Participants in both studies completed a comprehension check of three questions each. Any participant who failed one or more of these questions was immediately routed out of the experiment (16% in the money and 37% in the UberEats experiments). We additionally implemented an attention check at the end of the study, and exclude all participants who failed it (2% in the money and 1% in the UberEats experiments).

In the hypothetical money experiments, participants received $4.50 as a flat payment for completion of the study. In the UberEats study, participants received a completion fee of $4.00. In addition, one of the three parts of the experiment (intertemporal choice, risky choice, Raven IQ test) was randomly selected for payout, with associated probabilities of 25:5:70. Appendix G contains screenshots of all experimental instructions and comprehension checks.

4.5 Pre-Registration

Appendix Table 5 provides an overview of all treatments conducted for this paper, including pre-registration details. Our pre-registration includes (i) predictions 1–4 in Section 3, (ii) the prediction that cognitive uncertainty is correlated with across-trial choice variability, and (iii) descriptive analyses of the correlates of cognitive uncertainty to be discussed in Section 5.

[14]In our money experiments, average age is 42 years, 54% are female, and 45% have a college degree. In our UberEats experiments, average age is 28 years, 58% are female and 59% have a college degree.

[15]Because our experiments were conducted from late March through May 2021, we took various measures to ensure that only those prospective participants signed up for the study who were not concerned about ordering food for delivery due to COVID-19. First, the study description clarifies that people should not participate if they are concerned about ordering food for delivery due to COVID-19. Second, we restricted the sample to participants of age 45 and under. Third, we ask prospective participants whether they are worried about ordering delivery food due to COVID-19, and we immediately exclude anyone from the study whose response is affirmative. Finally, by late March 2021 it had become increasingly evident that delivery food is not a main source of COVID-19 transmission.

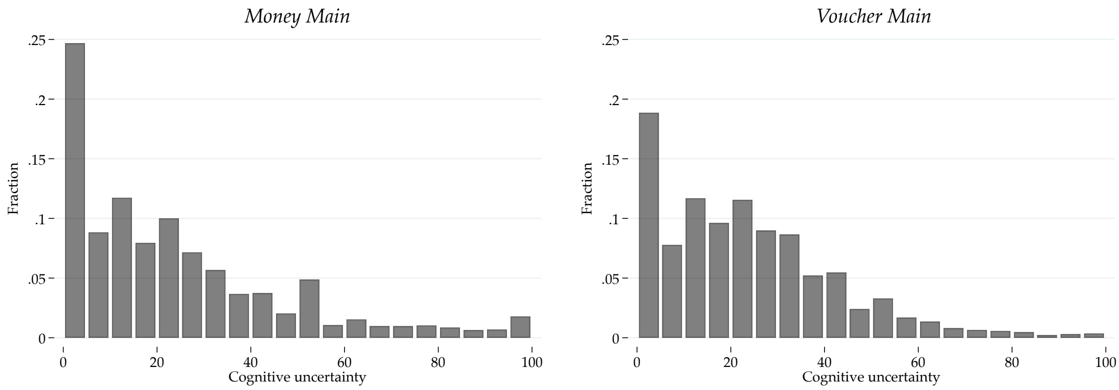

Figure 2: Histogram of cognitive uncertainty statements in *Money Main* (left panel, $N = 7,740$) and *Voucher Main* (right panel, $N = 6,000$).

5 Variation in Cognitive Uncertainty

5.1 Variation Across Participants and Decision Problems

Figure 2 shows histograms of task-level CU in the MPL decisions in treatments *Money Main* (left panel) and *Vouchers Main* (right panel), such that each participant corresponds to twelve observations. We see that 75% of all decisions in *Money Main* and 81% of decisions in *Voucher Main* are associated with strictly positive CU.

The heterogeneity in Figure 2 reflects both across-participant heterogeneity and systematic variation across choice problems. Figure 3 illustrates correlates of CU in treatment *Money Main* using binned scatter plots; the analogous figures for treatment *Voucher Main* look almost identical. The left panel shows that CU strongly increases in the length of the log time delay ($\rho = 0.16$), suggesting that payouts or consumption in two temporally distant periods are more difficult to evaluate against each other.

In light of this across-task variation, a relevant question is how stable people are in exhibiting high or low CU. In our data, participant-fixed effects explain 45-54% of the variation in CU. Thus, CU appears to have reasonably high within-domain stability. Looking at across-domain stability, the right panel of Figure 3 documents that a participant's average CU in intertemporal decisions is strongly correlated with the participant's average CU in separate risky choice (lottery) experiments that we implemented in the final part of our study. The raw correlation is $r = 0.62$ in *Money Main* and $r = 0.50$ in *Voucher Main*.[16]

[16]Other correlations between average subject-level CU and demographics are mostly small. The first value refers to the money study and the second one to the voucher study: $r = -0.08$ (0.01) with the score on Raven matrices IQ test, $r = -0.10$ (0.08) with age, $r = 0.06$ (0.06) with a female indicator, $r = -0.03$ (-0.05) with a college degree indicator, and $r = 0.07$ (-0.07) with log study completion time.

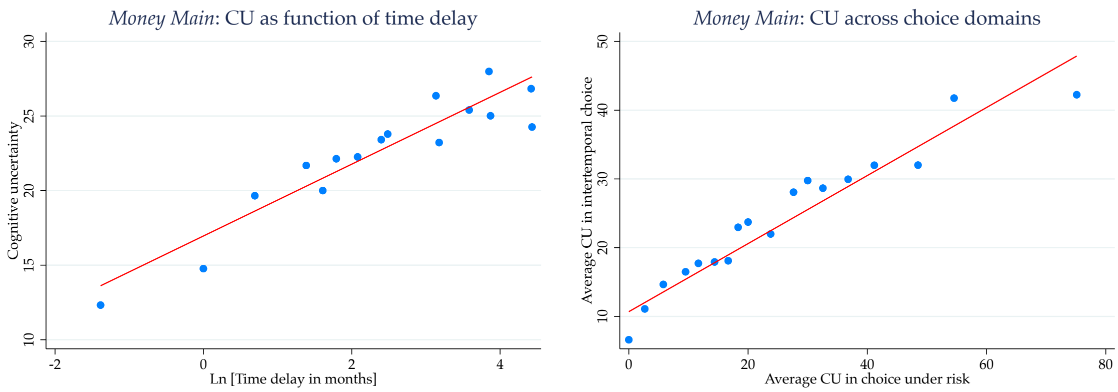

Figure 3: Binscatter plots. The left panel shows the relationship between task-level CU and the log time delay in a decision problem (N=7,740 decisions). The right panel shows the correlation between participant-level average CU in intertemporal choice and average CU in choice under risk (N=645 participants).

5.2 Is Cognitive Uncertainty Reflective of Actual Noise?

Some researchers have used choice variability as empirical measure of cognitive noise. We deem it useful to establish an empirical correspondence between our CU question and variability for two reasons. First, data on choice variability is useful to understand whether people's subjective perception of their own cognitive noise is roughly accurate. Second, a correlation between CU and choice variability may be seen as validation of a quantitative-but-unincentivized question, in the spirit of recent experimental validation studies in the literature (e.g. Falk et al., 2015; Enke et al., forthcoming).

Figure 4 shows the magnitude of across-trial variability as a function of cognitive uncertainty. Variability is computed as absolute difference in normalized switching points across two repetitions of the same choice list. We see that decisions that are associated with higher average CU across the two trials are more variable. In quantitative terms, an increase in average CU from zero to fifty is associated with a threefold increase in variability. In both datasets, the raw correlation is $\rho \approx 0.17$, $p < 0.01$.

Aside from choice variability, response times are sometimes used in the literature as an indicator for how difficult – and thus, potentially, cognitively noisy – choices are. We do not find a meaningful relationship between CU and response times in choice lists ($\rho = 0.03$), nor between response times and choice variability ($\rho = 0.04$).

6 Cognitive Uncertainty and Intertemporal Choice

6.1 Inelasticity of Decisions to the Time Delay

We begin by displaying the raw data: how intertemporal decisions vary as a function of the delay. For each choice list, a useful summary statistic is a participant's *normalized*

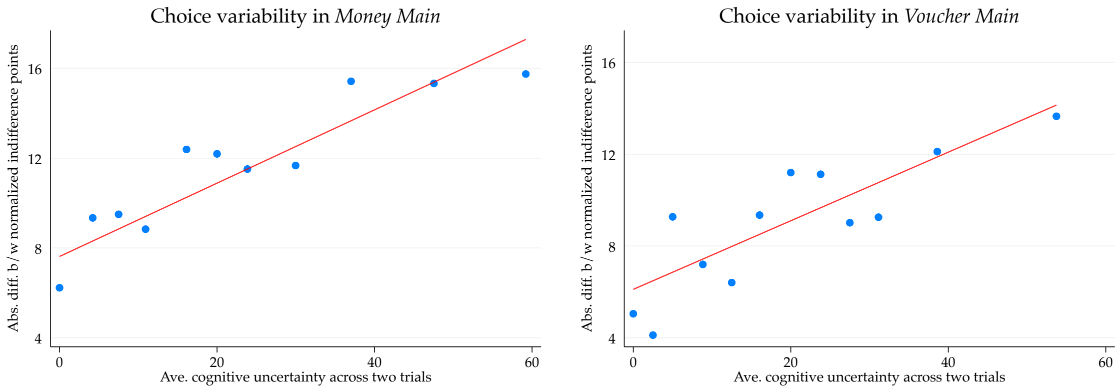

Figure 4: Link between cognitive uncertainty and across-task variability in normalized switch points in an exact repetition of the same decision problem in *Money Main* (left panel, $N = 1,290$) and *Voucher Main* (right panel, $N = 1,000$). The y-axis captures the absolute difference between the normalized indifference points across the two implementations. Average cognitive uncertainty is winsorized at 60 (roughly the 95th percentile in both datasets) for ease of visibility.

indifference point, which is given by the midpoint of the switching interval, divided by the later payment amount. This measure represents which payment at the earlier payment date makes the participant indifferent to receiving \$1 at the later date.

Figure 5 shows average normalized indifference points (in percent). The top panel shows results for treatment *Money Main* and the bottom panel those for *Voucher Main*. To make the results comparable, the x-axes are kept identical even though the maximal time delay in the vouchers study is only twelve months. We show results separately for participants with CU of zero and strictly positive CU.[17] For ease of illustration, we restrict attention to decision problems in which the early payment date is today, $t_1 = 0$. The analogous figure for $t_1 > 0$ looks very similar (Figure 12 in Appendix B).

The figure's main takeaway is that CU is strongly associated with compression of indifference points towards the center (roughly 50%). Notably, in treatment *Money Main*, this CU-associated inelasticity is sufficiently strong that cognitively uncertain participants act as if they are *less* patient over relatively short horizons, yet *more* patient over relatively long horizons, with a crossover point at around one year. This indicates that the main behavioral implication of CU in intertemporal choice is indeed insensitivity to time delays, rather than generically higher impatience. A second takeaway from Figure 5 is that behavior is very similar in *Money Main* and *Voucher Main*, including in its link to CU. In particular, cognitively uncertain decisions in *Voucher Main* also reflect lower patience over short time delays, yet as a result of insensitivity, this difference becomes ever smaller as the length of the time delay increases.

Table 1 presents corresponding OLS regression estimates. Here, we relate partici-

[17]This binary split is just for illustration purposes. Our regressions analyses leverage the entire variation in cognitive uncertainty.

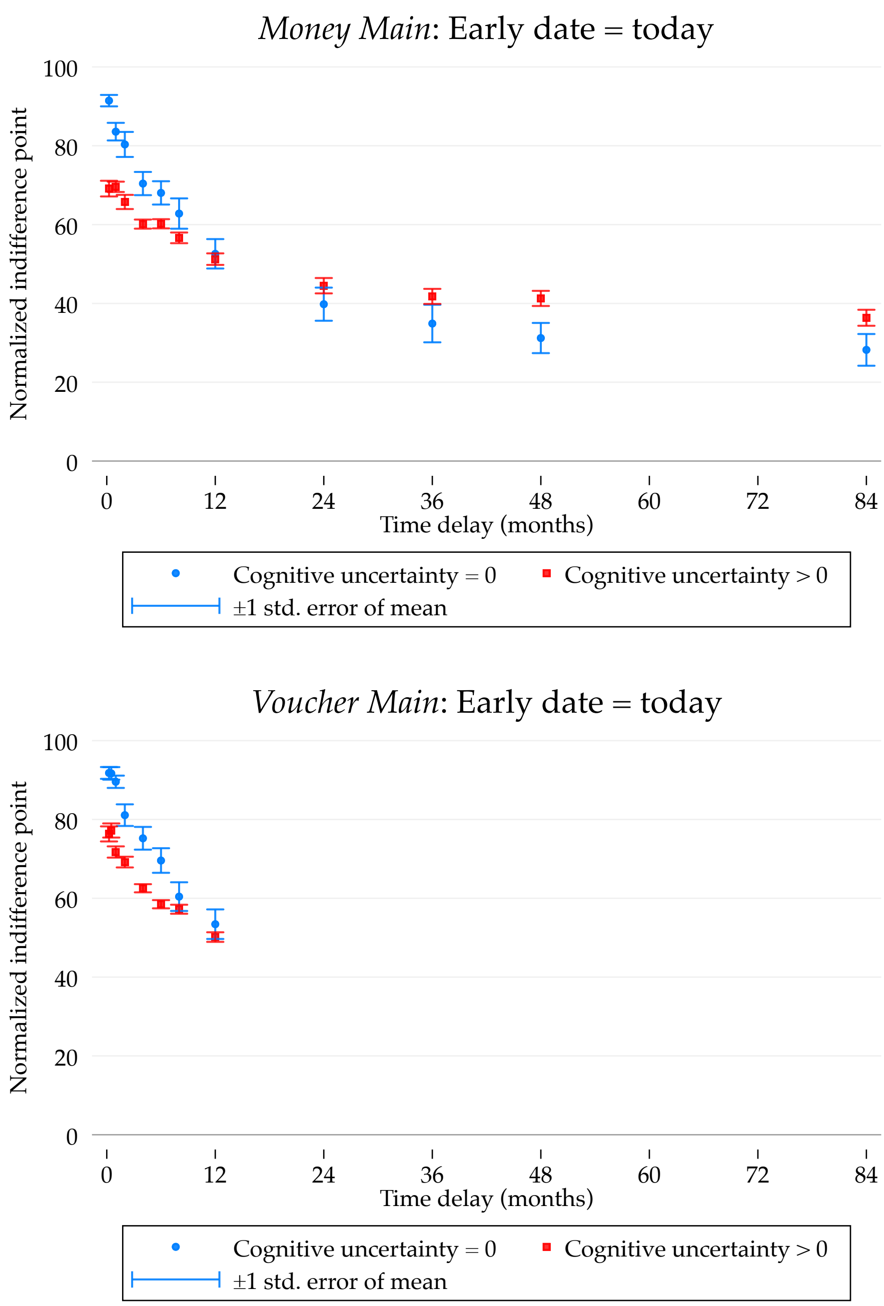

Figure 5: Observed discounting with $t_1 = 0$ in *Money Main* (top panel, $N = 4,948$) and *Voucher Main* (bottom panel, $N = 3,846$). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

pant's normalized indifference point to the length of the time delay, interacted with CU. Columns (1)–(4) show the results for *Money Main*, separately for whether the early payment date is today or in the future. Columns (5)–(8) show analogous results for *Voucher Main*. The results confirm the visual impression from Figure 5. First, CU is associated with a lower sensitivity of indifference values with respect to time delays, as can be in-

Table 1: Cognitive uncertainty and inelasticity with respect to time delays

	Dependent variable: Normalized indifference point							
Treatment:	*Money Main*				*Voucher Main*			
Sample:	$t1 = 0$		$t1 > 0$		$t1 = 0$		$t1 > 0$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Time delay (years)	-8.08***	-8.08***	-7.76***	-7.72***	-39.2***	-38.9***	-39.1***	-39.1***
	(0.39)	(0.39)	(0.39)	(0.39)	(2.16)	(2.14)	(3.88)	(3.86)
Time delay × Cognitive uncertainty	0.11***	0.11***	0.073***	0.071***	0.61***	0.59***	0.58***	0.59***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.08)	(0.08)	(0.14)	(0.14)
Cognitive uncertainty	-0.38***	-0.37***	-0.32***	-0.31***	-0.59***	-0.58***	-0.57***	-0.57***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.06)	(0.06)	(0.07)	(0.07)
Payment amount FE	No	Yes	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	4948	4948	2792	2792	3846	3846	2154	2154
R^2	0.17	0.19	0.19	0.21	0.20	0.21	0.13	0.14

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. Columns (1)–(4) include data from *Money Main*, where columns (1)–(2) restrict attention to decision problems with $t_1 = 0$ and columns (3)–(4) to problems with $t_1 > 0$. An analogous logic applies to columns (5)–(8) for *Voucher Main*. Demographic controls include age, gender and income bucket. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ferred from the positive interaction coefficient. Second, the regression intercept (which captures patience over very short horizons) is negatively correlated with CU, as we can infer from the significant raw CU term. These results are very similar for $t_1 = 0$ and $t_1 > 0$. A final comment regards the coefficient magnitudes. For example, in column (1), the coefficients suggest that increasing CU from zero to fifty (the 90th percentile) is associated with a decrease in sensitivity from 8.1 to 2.6 (or 68%), a large magnitude.[18]

6.2 Linking Cognitive Uncertainty to Empirical Regularities

6.2.1 Short-Run Impatience

Figure 5 provided strong visual evidence for the hypothesis that, over very short horizons, cognitively uncertain subjects are more impatient than cognitively certain ones, in both *Money Main* and *Voucher Main*. More formally, in *Money Main*, the raw correlation between normalized indifference points for one-week delays and cognitive uncertainty is $\rho = -0.45$ both when $t_1 = 0$ and when $t_1 > 0$. In *Voucher Main*, the same correla-

[18]The reason why the coefficient magnitudes are so different between *Money Main* and *Voucher Main* is the large difference in the average time delay between these two experiments. Once the data in *Money Main* are restricted to delays of at most one year, the coefficients are similar across the two experiments.

tions are given by $\rho = -0.39$ and $\rho = -0.45$. All of these correlations are statistically significant at the 1% level. Appendix Table 7 reports complementary regressions.

6.2.2 Decreasing Impatience

To study decreasing impatience, we follow the literature and define a *required rate of return* for a given normalized indifference point a^o as $RRR_{t_1,t_2}(a^o) \equiv \ln\left(\frac{c_{t_2}}{c_{t_1}}\right) = \ln\left(\frac{1}{a^o}\right)$. The RRR is a metric of impatience that depends on the delay. The literature frequently computes a per-period measure of patience as $RRR/\Delta t$. A transformation of this measure that captures per-period patience in an intuitive and structurally meaningful way is

$$\delta_H(a^o) \equiv e^{-RRR/\Delta t} = (a^o)^{1/\Delta t}. \tag{6}$$

This monotone transformation is attractive because – in a standard exponential discounting model without utility curvature and present bias – it directly corresponds to the exponential annual discount factor that is implied by the indifference point a^o. Thus, decreasing impatience says that $\delta_H(a^o)$ increases in the time delay, while under exponential discounting $\delta_H(a^o)$ is constant in the time delay.

Figure 6 shows the link between CU and decreasing impatience in four different panels: treatments *Money Main* and *Voucher Main*, separately for $t_1 = 0$ and $t_1 > 0$. Again, to make the figures comparable across experiments, we scale the x-axis to accommodate the longer time delays in *Money Main*. For each of the samples, we compute the average implied $\delta_H(a^o)$ across subjects for a given time horizon.[19]

The figures show that average per-period patience strongly increases in the time delay for cognitively uncertain participants. This is true in all four panels. For participants with CU of zero, however, per-period patience increases much more weakly. For example, for decisions in *Money Main*, implied per-period patience increases by a factor of 9.4 for choices associated with positive cognitive uncertainty (going from a time delay of one week to seven years), but by a factor of only 1.8 for decisions with zero cognitive uncertainty. Table 8 in Appendix C confirms these visual impressions through regressions.

The strong increase in per-period patience for high-CU decisions cannot be explained by present bias alone even if one asserted that CU and a desire for immediate gratification are correlated. This is because we find very similar patterns for $t_1 = 0$ and $t_1 > 0$,

[19]This visualization procedure is not subject to the aggregation insight of Weitzman (2001) and Jackson and Yariv (2014), which is that if the true data-generating process consists of subjects having different exponential discount functions, the average choice cannot necessarily be represented by an exponential function. This is not a problem here because we do not compute an implied δ_H for the average choice, but instead average the implied δ_H. Therefore, if the true process was exponential and participants had heterogeneous but constant discount factors, the average implied δ_H in Figure 6 should be constant in the time horizon. In any case, in our regression analyses, we always work with decision-level (rather than average) implied δ_H, which implies that potential aggregation issues never matter for our statistical tests.

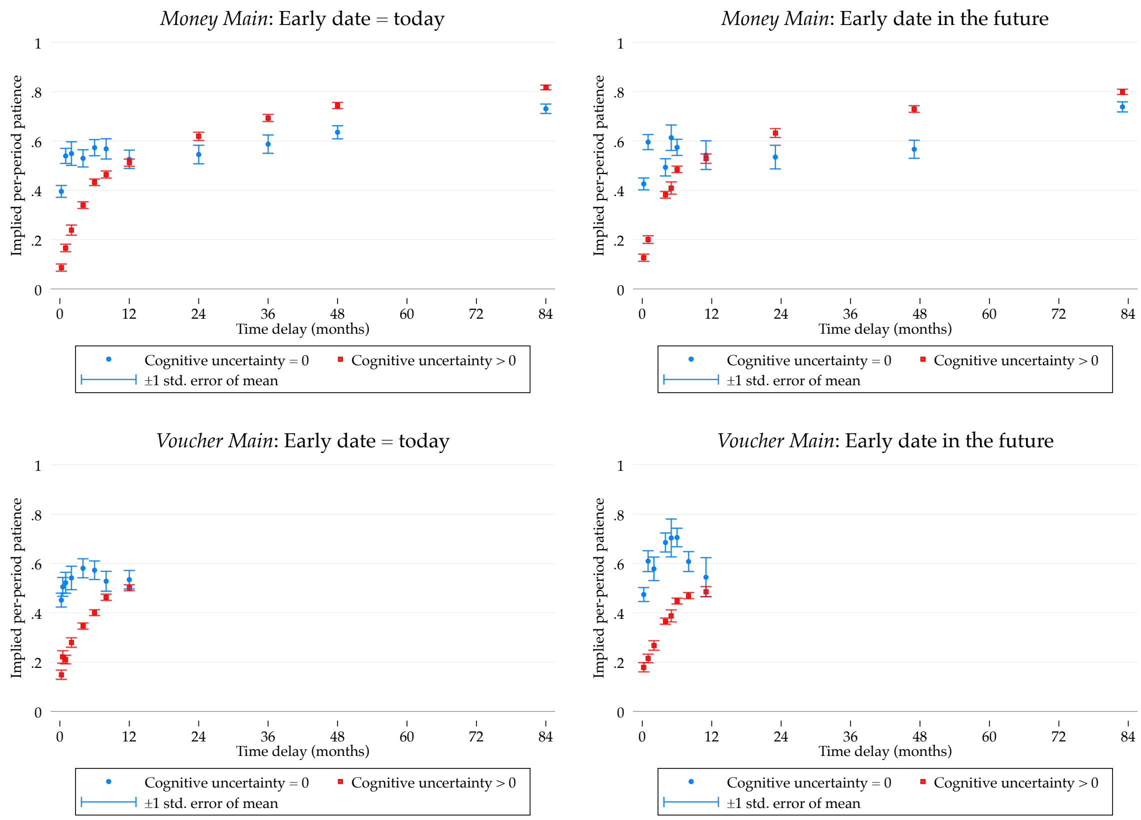

Figure 6: Implied per-period patience in *Money Main* (top panels) and *Voucher Main* (bottom panels), partitioned by whether the early payment date is today or in the future. Per-period patience is computed as $\delta_H(a^o) \equiv e^{-RRR/\Delta t} = (a^o)^{1/\Delta t}$, where a^o is the observed normalized indifference point. The figure shows average δ_H across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

while present bias only predicts diminishing impatience for $t_1 = 0$. Section 7 calibrates the relative importance of CU and present bias in generating observed behavior.

6.2.3 Subadditivity

We now turn to the two "subadditivity sets" in our data, each of which consists of three dates: set 1: {0, 6m, 12m}; set 2: {0, 4m, 8m}.[20] Following standard procedures in the literature, we compare the normalized indifference point obtained from the problem involving the one long interval with the product of the two normalized indifference points obtained from the respective two short intervals (the implied normalized indifference point of a long "composite interval"). Thus, although each subject makes three decisions for a given set, these give rise to two observations. Subadditivity occurs if the former quantity is larger than the latter. Table 2 summarizes the results for both *Money Main*

[20]Because we randomly selected some choice lists to be presented twice to the same participant, we sometimes have more than one observation for one of the three decisions that constitute a subadditivity set. In those cases, we average the decisions in the two identical choice lists.

Table 2: Cognitive uncertainty and subadditivity

	Dependent variable: Overall Normalized Indifference Point over Long Interval					
Treatment:	*Money Main*			*Voucher Main*		
	(1)	(2)	(3)	(4)	(5)	(6)
1 if one long interval, 0 if composite interval	8.53***	3.35**	3.63***	9.50***	1.51	1.56
	(0.62)	(1.32)	(1.32)	(0.60)	(1.61)	(1.60)
1 if one long interval × Cognitive uncertainty		0.25***	0.23***		0.32***	0.32***
		(0.06)	(0.06)		(0.06)	(0.06)
Cognitive uncertainty		-0.44***	-0.42***		-0.42***	-0.42***
		(0.06)	(0.06)		(0.08)	(0.08)
Set FE	Yes	Yes	Yes	Yes	Yes	Yes
Payment amount FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	No	No	Yes	No	No	Yes
Observations	1948	1948	1948	2000	2000	2000
R^2	0.02	0.07	0.09	0.05	0.08	0.09

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. Each subject makes three decisions for a given set, which give rise to two observations / composite normalized indifference points. The first is given by the normalized indifference point for a decision over the respective long horizon. The second is given by the product of the two normalized indifference points for the decisions over the two respective short horizons. Set fixed effects include fixed effect for each pair of decision problems that exhibit a front-end delay structure. Set 1: {0, 6m}, {6m, 12m} and {0m, 12m}. Set 2: {0, 4m}, {4m, 8m} {0m, 8m}. Demographic controls include age, gender and income. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

and *Voucher Main*. In both sets of experiments, we see strong evidence for the existence of subadditivity, see columns (1) and (4). In line with our hypothesis, the difference in observed patience between long and short intervals increases significantly in CU, see the interaction term in columns (2)–(3) and (5)–(6).

6.2.4 Front-End Delay Effects

Finally, we study the link between CU and front-end delay effects. These refer to the regularity that people exhibit greater patience in a decision problem in which both payment dates are moved forward by a constant. For example, people frequently appear more patient in tradeoffs between {6m, 12m} than between {0, 6m}. Recall that we predicted and pre-registered that cognitive uncertainty is uncorrelated with front-end delay effects. Therefore, an effective way to view these analyses is that they are a type of placebo exercise.

As summarized in Cohen et al. (2020), front-end delay effects are often but not always present in choices over monetary amounts. In our context, columns (1) and (4) document that we find highly significant and quantitatively large evidence for the presence of front-end delay effects. More importantly for our purposes, we find that the

Table 3: Cognitive uncertainty and front-end delay effects

	Dependent variable: Normalized indifference point					
Treatment:	*Money Main*			*Voucher Main*		
	(1)	(2)	(3)	(4)	(5)	(6)
1 if front end delay	3.07***	2.56*	2.47*	2.74***	4.98***	5.18***
	(0.85)	(1.32)	(1.30)	(0.86)	(1.68)	(1.67)
Front-end delay × Cognitive uncertainty		0.048	0.049		-0.055	-0.064
		(0.05)	(0.05)		(0.05)	(0.05)
Cognitive uncertainty		-0.30***	-0.28***		-0.24***	-0.23***
		(0.05)	(0.05)		(0.06)	(0.06)
Set FE	Yes	Yes	Yes	Yes	Yes	Yes
Payment amount FE	No	No	Yes	No	No	Yes
Demographic controls	No	No	Yes	No	No	Yes
Observations	2393	2393	2393	2337	2337	2337
R^2	0.00	0.05	0.06	0.01	0.05	0.05

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. Set fixed effects include fixed effect for each pair of decision problems that exhibit a front-end delay structure. Set 1: {0, 6m} and {6m, 12m}. Set 2: {0, 4m} and {4m, 8m}. Demographic controls include age, gender and income. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

correlation between front-end delay effects and cognitive uncertainty is either small and statistically insignificant (columns (2)–(3)) or even goes in the opposite direction (columns (5)–(6)). This is despite a relatively large sample size of $N = 2,393$ decisions (645 subjects) in *Money Main* and $N = 2,337$ decisions (500 subjects) in *Voucher Main*.

6.3 Taking Stock: Modeling Approaches vs. Evidence

As noted earlier, our primary contribution is to document the relevance of noisy cognition for intertemporal choice, rather than to definitively disentangle different classes of random choice models that often make similar predictions (and each of which come in different variants). This being said, a comparison of the empirical results with the discussion in Section 3 allows us to draw some tentative conclusions about which types of models explain the patterns better than others. A crucial role in this regard play the choice patterns regarding subadditivity and front-end delay effects. The main reason is that the cognitive-noise-in-action-space framework and the random response model that we sketched in Section 3 predict that cognitive noise is correlated with subadditivity but not with front-end delay effects. Random preference models and the cognitive noise model of Gabaix and Laibson (2017), on the other hand, both predict that cognitive noise is linked to front-end delay effects but not to subadditivity. Given that we find that cognitive uncertainty is predictive of subadditivity but not of front-end delay effects, we conclude that random preference models and the approach of Gabaix and Laibson (2017) and Gershman and Bhui (2019) do not explain all aspects of the evidence.

6.4 Robustness

Omitted variables. Given that all analyses up to this point are correlational in nature, a potential concern is the existence of a stable participant characteristic other than cognitive uncertainty that somehow generates the results. While we are not aware of other characteristics that could plausibly lead to higher implied impatience over short horizons, yet lower implied impatience over long horizons, we perform a robustness check by including participant fixed effects in our main regression in Table 1. By definition, these soak up fixed subject characteristics such as overall cognitive ability. As a result, inelasticity with respect to variation in the time delay is identified purely off of within-participant-across-task variation in CU. As we document in Appendix Table 9, the results remain statistically significant conditional on these subject fixed effects.

Direct elicitation experiments. Up to this point, all results were derived from experiments in which intertemporal choice behavior was elicited using choice lists. To document that the logic of CU and inelasticity extends to another elicitation technique, treatment *Money Main* also included a direct elicitation component, see Section 4. Here, subjects were directly asked how much they value a hypothetical payment of \$y in $t = t_2$ in terms of a payment to be received today. To answer this question, subjects directly typed a dollar amount into a text box. After each decision, subjects indicate their cognitive uncertainty by indicating their subjective probability that their true valuation for the later payment actually lies within ±\$1 of their stated valuation.

Appendix D shows that these direct elicitation experiments deliver very similar results as the ones reported thus far. Specifically, we find that (i) CU is significantly correlated with across-trial choice variability; (ii) CU is strongly correlated with short-run impatience over one week; (iii) CU is correlated with decreasing impatience; (iv) CU is correlated with subadditivity; and (iv) CU is again uncorrelated with front-end delay effects. Thus, all of our results from the MPL experiments replicate using the direct elicitation technique.

7 Model Estimations

We proceed by estimating eq. (2) from Section 3 to gauge how well such a reduced-form model fits the data, and how much the measurement of cognitive uncertainty contributes to model fit. In eq. (2), the weight λ depends on the magnitude of cognitive noise. We do not observe cognitive noise itself but cognitive uncertainty, denoted p_{CU}. We proceed by using the heuristic approximation $\lambda = 1 - \alpha p_{CU}$, where $\alpha \geq 0$ is a nuisance parameter to be estimated. With CRRA utility and larger-later payment $x \equiv c_{t_2} \geq 1$, eq. (2) suggests that the mean observed choice in our experiments is determined as

$$
\begin{aligned}
\mathbb{E}[a^o] &= \lambda(p_{CU}) \cdot \mathbb{E}[s] \cdot x + (1 - \lambda(p_{CU})) \cdot d \cdot x \\
&= (1 - \alpha \cdot p_{CU}) \cdot (\delta^{\Delta t})^{1/\gamma} \cdot x + (\alpha \cdot p_{CU}) \cdot d \cdot x
\end{aligned}
\tag{7}
$$

This equation, amended by a mean-zero error term, can be estimated using straightforward nonlinear least squares techniques. Specifically, we observe a^o, Δt and p_{CU}, and estimate δ, d and α.[21] To assess and compare model fit, we estimate four model variants. First, a baseline exponential discounting model in which we ignore CU (i.e., we set $\alpha = 0$). Second, also for benchmarking purposes, a $\beta - \delta$ model, which also precludes a role for CU. Finally, we estimate both of these variants including CU.[22]

Notice that the estimate of d has two potential interpretations. Under the Bayesian cognitive noise interpretation, d is a constant cognitive default action that people anchor on. Under the random response interpretation, d is the mean of the distribution function $F(\cdot)$ from which random responses are drawn.

Aggregate estimates. We begin by estimating the model across subjects, treating the data as if it was generated by one representative agent. Table 4 summarizes the model estimates across the three different types of experiments that we report in this paper. There are five main takeaways. First, in line with prior research, a pure exponential discounting model fits the data poorly. Second, a beta-delta model fits the data considerably better, but not nearly as well as a model that includes both exponential discounting and CU (see the Akaike Information Criterion values in the last row). Third, a model that includes both a role for taste-based present bias and CU performs best. This – in line with our results on front-end delay effects – again highlights that a desire for immediate gratification and cognitive noise are distinct and complementary objects. Fourth, ignoring CU in the estimations considerably inflates the role of present bias β. Fifth, the estimates are strikingly similar across experiments; in particular, the estimated d is always around 50% of the larger-later reward.

To put these estimates in perspective, note that our setup in which the earliest reward lies several hours in the future likely underestimates the true extent of present bias. Even

[21]The risk aversion parameter, γ, is separately estimated on our risky choice experiments in the final part of the study, and taken as given in the intertemporal choice estimations.

[22]The amended estimation equation for $\beta - \delta - CU$ is given by $a^o = (1 - \alpha \cdot p_{CU}) \cdot (\beta \delta^{\Delta t})^{1/\gamma} + (\alpha \cdot p_{CU}) \cdot d$.

Table 4: Estimates of model parameters across experiments

	Money Main MPL				*Money Main Direct Elicitation*				*Voucher Main MPL*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	δ	$\beta-\delta$	$\delta-CU$	$\beta-\delta$ $-CU$	δ	$\beta-\delta$	$\delta-CU$	$\beta-\delta$ $-CU$	δ	$\beta-\delta$	$\delta-CU$	$\beta-\delta$ $-CU$
$\hat{\delta}$	0.96	0.98	0.97	0.98	0.97	0.99	0.98	0.99	0.94	0.95	0.95	0.95
$\hat{\beta}$		0.77		0.86		0.76		0.85		0.89		0.95
$\hat{d}$			0.51	0.49			0.52	0.49			0.57	0.56
AIC	64,148	63,165	61,904	61,701	32,247	31,391	30,993	30,791	46,980	46,652	45,853	45,817

Notes. Estimates of different versions of (7). MPL = multiple price list. AIC = Akaike Information Criterion. Each column corresponds to a separate model estimation. Columns (1), (5), (9): set $\beta = 1$ and $\alpha = 0$. Columns (2), (6), (10): set $\alpha = 0$. Columns (3), (7), (11): set $\beta = 1$. All estimated standard errors (computed based on clustering at the subject level) are smaller than 0.02. In estimations that include CU, we also estimate the nuisance parameter α (not reported). All estimations are conducted by setting a CRRA parameter of $\gamma = 0.94$, which is the population-level risk aversion that was separately estimated on the risky choice data. The exponential parameter δ is the monthly discount factor.

though we find clear evidence for $\beta < 1$, recent experimental work suggests that most discounting occurs in the first few hours following a decision (e.g., Augenblick, 2018), something that is not captured in our experimental paradigm.

Figure 7 visualizes the fit of the various estimated models for treatment *Money Main*, separately for decision problems in which the early payment date is today or in the future.[23] The figures are constructed by generating predicted values, based on the parameter estimates in Table 4. We again see that exponential discounting fits the data very poorly. Likewise, almost by construction, the canonical beta-delta model fits poorly when the early payment date is in the future.[24] On the other hand, when the early payment date is today, the beta-delta model performs well in fitting behavior over relatively short time delays, but (as is well-known) performs relatively poorly in capturing the strong flattening out of the observed data for long time delays.

The delta-CU model, on the other hand, captures several key aspects of the data. First, it partly accounts for some of the extreme impatience over short horizons. Second, the model accounts much better (though also somewhat imperfectly) for the strong compression effects over long horizons. Third, the delta-CU model matches the data reasonably well both when the early payment date is today and when it is in the future.

Individual-Level Estimates. Estimating any intertemporal choice model at an aggregate level is problematic because participants might have heterogeneous discount factors (Weitzman, 2001; Jackson and Yariv, 2014). Therefore, we proceed by estimating the

[23]Appendix Figure 14 shows analogous results for *Voucher Main*.

[24]The different model fit for the exponential discounting and the beta-delta model for the case of $t_1 > 0$ result from the fact that we estimate both models on *all* data, including those with $t_1 = 0$.

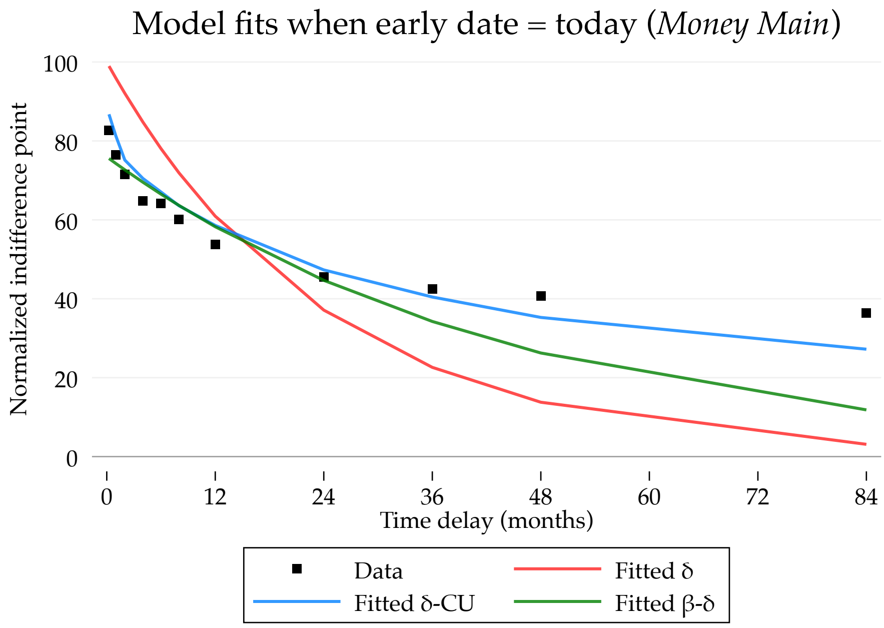

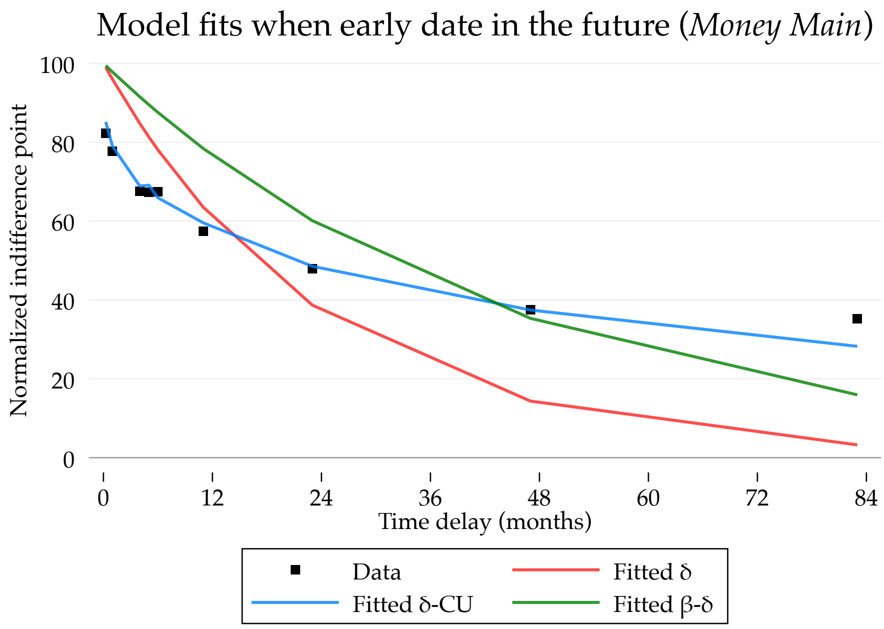

Figure 7: Model fit vs. data in *Money Main*. The model predictions are computed as fitted values of the parameter estimates in Table 4.

model at the level of individual subjects.[25] We report the results in Appendix Table 10. To summarize, there is substantial individual-level variation in estimated model parameters. For most parameters, the center of the estimated coefficient distributions is well in line with the parameters in our representative-agent estimation.[26] The empirical distri-

[25] To increase power in these individual-level estimations, we restrict attention to treatment *Money Main*, in which each subject completed both 12 MPLs and 6 direct elicitation tasks.

[26] An exception is the the present bias parameter β. We find *less* pronounced present bias (larger β) in our individual estimations than the aggregate ones, in line with the theoretical insight that aggregate

bution of $\hat{d}$ is roughly bell-shaped with a center around $\hat{d} \approx 0.5$, with two pronounced spikes at $\hat{d} = 1$ and $\hat{d} = 0$.

Discussion. Our estimations consistently suggest that a potential cognitive default action or mean random response is given by roughly 50% of the larger-later payment. This is true both in experiments with price lists and in direct elicitation experiments. Of course, given the available evidence, we do not intend to take a strong stance on whether this estimate will be context-specific. While we suspect that it will be (see the discussion in the Conclusion), it is also interesting to note that the "central" nature of the estimated d jives well with a large body of work in both economics and psychology that suggests that people's heuristic responses to decision problems tend to be *intermediate* in nature. In psychology, this well-known finding has come to be known as the "central tendency effect" (Hollingworth, 1910), which appears across a large set of decision domains. Indeed, in joint work with cognitive psychologists, we have documented that central tendency effects in various perceptual domains are strongly linked to cognitive uncertainty (Xiang et al., 2021). In economics, a related effect is the so-called compromise effect (see, e.g., Beauchamp et al., 2019, for an example in risky choice), which captures that people tend to indicate indifference values that are in the "center" of a choice set.

8 Why Cognitive Microfoundations Matter

It may not be immediately obvious why it is important for economists to understand that intertemporal choice is to a large degree governed by bounded rationality rather than only preferences, if reduced-form discount functions such as variants of the generalized hyperbola generally perform reasonably well in fitting data (even if for the wrong reasons). While we believe that understanding cognitive microfoundations is scientifically valuable in its own right, we now additionally show that understanding cognitive mechanisms matters for economic predictions and, potentially, policy.

8.1 Complexity and Hyperbolic Discounting

A main implication of a preferences-based account is that the hyperbolic shape of discounting is fixed. Our account, on the other hand, predicts that economically-relevant phenomena such as short-run impatience and hyperbolic discounting will be more pronounced in environments that increase cognitive noisiness. In the absence of theoretical

quasi-hyperbolic discounting can partly result from the aggregation of preferences of individuals with different discount factors.

guidance for what determines cognitive noise, we conjecture that the magnitude of cognitive noise will be a function of (i) the complexity of the decision problem and (ii) the availability of cognitive resources for deliberation of the problem.

Task complexity. Because there is no widely accepted definition for task complexity, we implement two treatments that plausibly increase the perceived complexity of the intertemporal decision problems. Here, one treatment is aimed at increasing the complexity of the time delay, while the other treatment increases the complexity of mentally simulating payoffs and resulting utils. Specifically, in *Money Complex Dates*, we implemented the same procedures as in *Money Main*, except that all payout dates in the choice lists were represented as a math task. For instance, "In 1 year" could be represented as "In (6*2/3-3) years AND (3*6/2-9) months AND (5*4/2-10) days." Appendix Figure 18 provides a screenshot.

In treatment *Money Complex Amounts*, we again implemented the same procedures as in *Money Main*, except that, for the delayed option A in a choice list, the monetary amount was again represented as a math problem, such as "$(4*8/2)+(8*9/2)-12". Appendix Figure 19 provides a screenshot. Relative to our baseline condition, this treatment leaves the complexity of the display of payment dates constant, but makes determining the consumption implications of a choice option more difficult.

Availability of cognitive resources. To manipulate the cognitive resources that people have at their disposal to mentally simulate their true indifference point, prior literature has worked with cognitive load or time pressure / waiting periods (Deck and Jahedi, 2015; Imas et al., 2021; Ebert, 2001). The main result in this literature is that having fewer cognitive resources available generally leads to lower revealed patience. Yet, in these studies, researchers implemented relative short time delays. Our account predicts that people act as if they are actually *more* patient over long horizons when they have fewer cognitive resources available. In our treatment *Money Load*, participants are tasked with simultaneously (i) completing the intertemporal choice problems over money described in the previous section and (ii) adding up red numbers that appeared at random intervals next to the choice list.[27]

An obvious issue with our complexity and load manipulations is that cognitive effort and resulting response times are endogenous: in principle, it is conceivable that subjects in the more complex conditions take substantially longer to complete the tasks, so that no effect on CU would be visible. To prevent this, we implemented a time limit of 25 seconds per choice list in each of these conditions, including in a replication of treatment

[27]We also separately implemented a within-subject design that manipulated the presence of the number counting task within subjects across tasks. The results are very similar, see Appendix F.

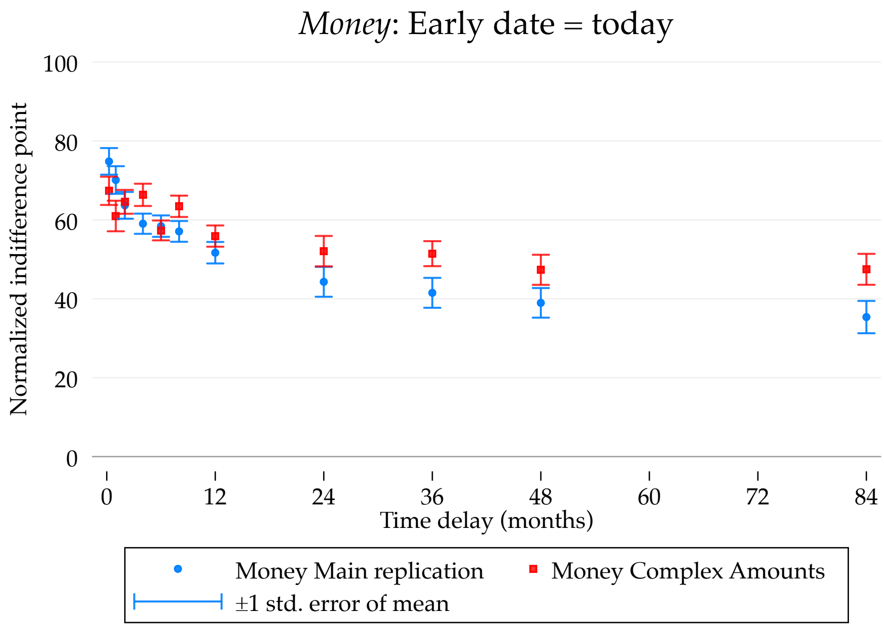

Figure 8: Observed discounting with $t_1 = 0$ in *Money Main replication* ($N = 1,932$) and *Money Complex Amounts* ($N = 1,836$). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

Money Main that we administered in the same experimental sessions. See Appendix E for example screenshots for all treatments. We conducted these experiments with a separate sample of 617 participants, in which each participant was randomly assigned to one of the four treatments (*Money Complex Dates*, *Money Complex Amounts*, *Money Load* and *Money Main Replication*).

Results. We find that all three treatment variations substantially increase stated CU relative to the replication of our main treatment. The magnitude of the increase is between 5 and 12 percentage points (20% to 50%, respectively), $p < 0.01$ for all comparisons to the baseline replication. Turning to intertemporal choice behavior, Figure 8 summarizes the results for treatment *Money Complex Amounts*. We see that the indifference points in the more complex treatment are more compressed around 50% compared to the replication of the baseline treatment. As a result, participants in the more complex treatment behave as if they are more impatient over short horizons but less impatient over long ones. Appendix E summarizes the results for treatments *Money Complex Dates* and *Money Load*, which look very similar. See Appendix Table 17 for statistical tests. In all, we see that acknowledging a role of cognitive noise facilitates improved predictions about the context-dependence of hyperbolicity.

8.2 Choice Architecture

In contrast to preferences-based theories of extreme short-run impatience, an account of cognitive uncertainty predicts that short-run impatient choices will often be associated with a sense of "nervousness" that the decision reflects an error. Thus, people may be open to (or even actively seek out) advice about how to behave. To study the relevance of cognitive uncertainty for choice architecture, we test whether it is indeed true that people with cognitive uncertainty about a given decision are more likely to follow the advice of an outside expert. This is arguably a strong hypothesis because variation in intertemporal decisions surely partly reflects genuine heterogeneity in preferences (e.g., in δ). Given that outside experts will rarely know the decision-maker's true preferences, following the advice of an expert is a double-edged sword: it may reduce the probability of making mistakes, but increase the probability of doing something that goes against one's individual preferences.

To assess the relevance of cognitive uncertainty for advice-seeking and choice architecture, we implement treatment *Voucher Advice*. This treatment follows exactly the same protocol as *Voucher Main*, except that it introduces a piece of advice. In the first choice list, we fixed the early payment date at today and varied the delayed payment date between one week and two months. After the participant had indicated their decisions in this choice list and their cognitive uncertainty, we presented a surprise announcement:[28]

> *We surveyed a few academic economists about which advice they would give to participants in this study regarding which decisions to make. These academic economists recommend that participants choose the delayed Voucher A in all rows of the choice list you just completed. We recognize that decisions like these depend on your own preferences, so we neither encourage nor discourage you to follow this advice. However, should you wish to revise your decision, you can do so in the choice list below. The choices that are indicated right now are those that you made yourself a few seconds ago.*

We pre-registered the sample size and our prediction that cognitive uncertainty is associated with a higher likelihood of following expert advice by revising a previous decision at `https://aspredicted.org/jk5s5.pdf`.

Experimentals like these are potentially subject to experimenter demand effects, according to which participants revise their decisions purely because they believe that the

[28]No deception was involved in the design of the study because we actually polled Harvard-based economists for advice. We suspect that the reason why people are comfortable articulating advice in such situations is that – over timeframes of one week to two months as in our study – even mildly impatient decisions imply absurdly high discount rates.

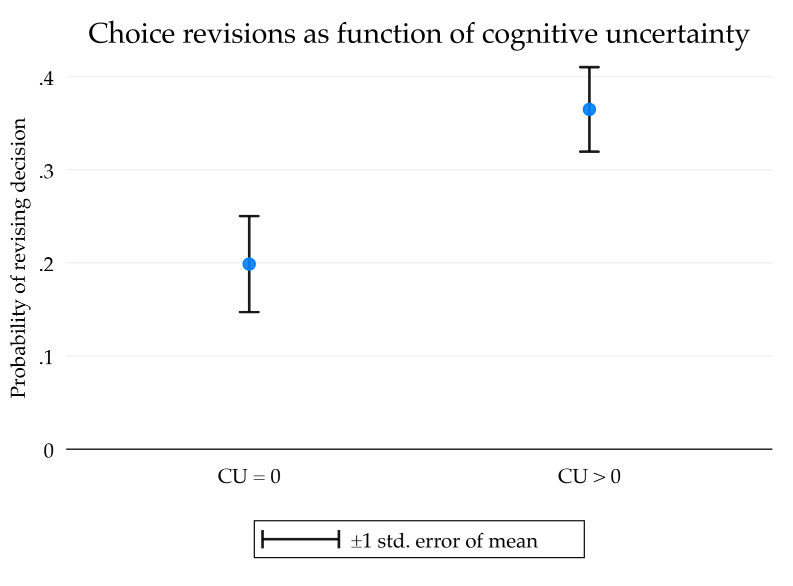

Figure 9: Probability of revising decision towards higher patience, as a function of cognitive uncertainty ($N = 153$). The figure is constructed controlling for the normalized indifference point before seeing advice. In other words, the y-axis shows the residual probability of revising the decision after the initial choice is partialed out through an OLS regression.

experimenter would like them to. In our context, this "level effect" is irrelevant because we are only interested in the differential responsiveness to advice of participants with and without cognitive uncertainty. Our identifying assumption is therefore that cognitively uncertain subjects are not subject to stronger demand effects.

In our data, 34% of participants revise their decision upon seeing advice, where almost all revisions are in the direction of higher patience. Figure 9 shows the relationship between cognitive uncertainty and choice revisions towards the advice of full patience.[29] We see that participants with strictly positive cognitive uncertainty are 16 percentage points (80%) more likely to revise their choice, $p < 0.01$.

9 Discussion

Contribution. Much of behavioral economics views intertemporal choice, and famous empirical regularities, as largely determined by non-standard discount functions (preferences). This paper argues for and empirically documents an important role of cognitive noise and complexity for intertemporal decision-making. An innovation of our study is that we directly measure and exogenously manipulate cognitive noise through self-reported cognitive uncertainty. Using this tool, we document that a large share of

[29]Because subjects with higher cognitive uncertainty on average state lower indifference points in their initial decision, they have more "room" to adjust. We control for this by residualizing the y-axis of Figure 9 from the initial normalized indifference point through a linear regression (the results are even stronger without this adjustment).

short-run impatience and hyperbolic discounting are driven by bounded rationality and cognitive noise, rather than impatient preferences. These insights matter not just from a scientific perspective but arguably have real economic implications. On the intensive margin of decision-making, hyperbolic discounting depends on complexity. On the extensive margin, cognitively uncertain people welcome the advice of experts even when those experts don't know their preferences. While we emphasize throughout the paper that cognitive noise is complementary to (rather than replaces) taste-based present bias, we have shown that cognitive noise provides a better account of many of key economic phenomena that are often ascribed to present bias. In all, we interpret these results as providing some of the first direct empirical evidence that cognitive noise and cognitive uncertainty are relevant for a broad set of economic aspects of intertemporal choice.

Link to cognitive effects in intertemporal choice research. We conjecture that our account of cognitive uncertainty provides a rationale for extant empirical findings about "cognitive" effects in intertemporal choice research. The perhaps most widely-known result on cognition and intertemporal choice is that, if the time delay is relatively short, a lower availability of cognitive resources is associated with less patient decisions. At the same time, Ebert (2001) presents evidence that suggests that, over long horizons, a lower availability of resources makes people *more* patient. Our account of the link between inelasticity and cognitive uncertainty reconciles this somewhat puzzling combination of results. Moreover, Cubitt et al. (2018) present intriguing evidence that people's decisions are much less sensitive to variation in the time delay when intertemporal decisions involve cross-domain comparisons (car now vs. vacation later) than when they only concern within-domain comparisons (car now vs. nicer car later). While no preferences-based intertemporal model predicts such effects, we conjecture that they are driven by higher cognitive noisiness in cross-domain comparisons.

Limitations. Our paper does not purport to explain nearly all intertemporal choice anomalies. One regularity that our study does not address are well-known framing effects, such as the speed-up / delay asymmetry (Loewenstein and Prelec, 1992) or date / delay effects (Read et al., 2005). At the same time, we do conjecture a potential link between such framing effects and our work: if one choice option is presented to people as the default that they can "speed up" at a cost, it seems plausible that people use that option as a cognitive default. Based on this idea, we conjecture that speed-up / delay asymmetries are more pronounced when cognitive uncertainty is high.

More generally, this conjecture highlights that further research is needed to understand potential cognitive default actions. In this paper, we estimate the default action to be "intermediate," which is consistent with various documentations of central tendency effects in cognitive psychology. Yet, it is important to note that the specific intertemporal choice context we study is one with which people have little or no experience. Future research will explore how potential cognitive default actions depend on experience and contextual influences.

References

Abeler, Johannes and Simon Jäger, "Complex Tax Incentives," *American Economic Journal: Economic Policy*, 2015, *7* (3), 1–28.

Agranov, Marina and Pietro Ortoleva, "Stochastic choice and preferences for randomization," *Journal of Political Economy*, 2017, *125* (1), 40–68.

__ **and** __ , "Ranges of Preferences and Randomization," *Working Paper*, 2020.

Andersen, Steffen, Glenn W Harrison, Morten I Lau, and E Elisabet Rutström, "Discounting behavior: A reconsideration," *European Economic Review*, 2014, *71*, 15–33.

Andreoni, James, Christina Gravert, Michael A Kuhn, Silvia Saccardo, and Yang Yang, "Arbitrage or narrow bracketing? On using money to measure intertemporal preferences," Technical Report, National Bureau of Economic Research 2018.

Beauchamp, Jonathan P, David I Laibson, and Christopher F Chabris, "Measuring and controlling for the compromise effect when estimating risk preference parameters," *Experimental Economics*, 2019, pp. 1–31.

Benzion, Uri, Amnon Rapoport, and Joseph Yagil, "Discount rates inferred from decisions: An experimental study," *Management science*, 1989, *35* (3), 270–284.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Memory, attention, and choice," *The Quarterly journal of economics*, 2020.

Brocas, Isabelle, Juan D Carrillo, and Jorge Tarrasó, "How long is a minute?," *Games and Economic Behavior*, 2018, *111*, 305–322.

Bulley, Adam, Karolina Lempert, Colin Conwell, and Muireann Irish, "Intertemporal choice reflects value comparison rather than self-control: insights from confidence judgments," *PsyArXiv*, 2021.

Butler, David J and Graham C Loomes, "Imprecision as an account of the preference reversal phenomenon," *American Economic Review*, 2007, *97* (1), 277–297.

Carrera, Mariana, Heather Royer, Mark Stehr, Justin Sydnor, and Dmitry Taubinsky, "Who Chooses Commitment? Evidence and Welfare Implications," Technical Report, National Bureau of Economic Research 2019.

Cerreia-Vioglio, Simone, David Dillenberger, and Pietro Ortoleva, "Cautious expected utility and the certainty effect," *Econometrica*, 2015, *83* (2), 693–728.

Chakraborty, Anujit, "Present bias," *Econometrica*, 2020.

— **, Yoram Halevy, and Kota Saito**, "The relation between behavior under risk and over time," *American Economic Review: Insights*, 2020, *2* (1), 1–16.

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White, "Measuring time preferences," *Journal of Economic Literature*, 2020, *58* (2), 299–347.

Cubitt, Robin, Rebecca McDonald, and Daniel Read, "Time matters less when outcomes differ: Unimodal vs. cross-modal comparisons in intertemporal choice," *Management Science*, 2018, *64* (2), 873–887.

Curry, David, "Uber Eats Revenue and Usage Statistics (2021)," 2021.

Dasgupta, Partha and Eric Maskin, "Uncertainty and hyperbolic discounting," *American Economic Review*, 2005, *95* (4), 1290–1299.

Deck, Cary and Salar Jahedi, "The effect of cognitive load on economic decision making: A survey and new experiments," *European Economic Review*, 2015, *78*, 97–119.

Dertwinkel-Kalt, Markus, Holger Gerhardt, Gerhard Riener, Frederik Schwerter, and Louis Strang, "Concentration bias in intertemporal choice," *Review of Economic Studies*, 2021.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde, "Are risk aversion and impatience related to cognitive ability?," *American Economic Review*, 2010, *100* (3), 1238–60.

Ebert, Jane EJ, "The role of cognitive resources in the valuation of near and far future events," *Acta psychologica*, 2001, *108* (2), 155–171.

— **and Drazen Prelec**, "The fragility of time: Time-insensitivity and valuation of the near and far future," *Management science*, 2007, *53* (9), 1423–1438.

Enke, Benjamin and Thomas Graeber, "Cognitive uncertainty," Technical Report, National Bureau of Economic Research 2020.

— **, Ricardo Rodriguez-Padilla, and Florian Zimmermann**, "Moral Universalism: Measurement and Economic Relevance," *Management Science*, forthcoming.

Epper, Thomas, Ernst Fehr, Helga Fehr-Duda, Claus Thustrup Kreiner, David Dreyer Lassen, Søren Leth-Petersen, and Gregers Nytoft Rasmussen, "Time discounting and wealth inequality," *American Economic Review*, 2020, *110* (4), 1177–1205.

Ericson, Keith Marzilli and David Laibson, "Intertemporal choice," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 2, Elsevier, 2019, pp. 1–67.

Esponda, Ignacio and Emanuel Vespa, "Hypothetical Thinking: Revisiting Classic Anomalies in the Laboratory," *Working Paper*, 2016.

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde, "The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences," *Working Paper*, 2015.

Fennell, John and Roland Baddeley, "Uncertainty plus prior equals rational bias: An intuitive Bayesian probability weighting function.," *Psychological Review*, 2012, *119* (4), 878.

Frydman, Cary and Lawrence Jin, "Efficient Coding and Risky Choice," *Quarterly Journal of Economics*, 2021.

_ **and Salvatore Nunnari**, "Cognitive Imprecision and Strategic Behavior," *Working Paper*, 2021.

Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki, "Speed, accuracy, and the optimal timing of choices," *American Economic Review*, 2018, *108* (12), 3651–84.

Gabaix, Xavier, "Behavioral inattention," in Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics-Foundations and Applications 2*, Elsevier, 2019, p. 261.

_ **and David Laibson**, "Myopia and discounting," Technical Report, National bureau of economic research 2017.

Gershman, Samuel J and Rahul Bhui, "Rationally inattentive intertemporal choice," *bioRxiv*, 2019, p. 680652.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel, "Very long-run discount rates," *The Quarterly Journal of Economics*, 2015, *130* (1), 1–53.

Gupta, Neeraja, Luca Rigott, and Alistair Wilson, "The Experimenters' Dilemma: Inferential Preferences over Populations," *arXiv preprint arXiv:2107.05064*, 2021.

Halevy, Yoram, "Strotz meets Allais: Diminishing impatience and the certainty effect," *American Economic Review*, 2008, *98* (3), 1145–62.

__ , "Some comments on the use of monetary and primary rewards in the measurement of time preferences," *Unpublished manuscript*, 2014.

Harrison, Glenn W, Morten I Lau, and E Elisabet Rutström, "Dynamic consistency in Denmark: A longitudinal field experiment," 2005.

He, Lisheng, Russell Golman, and Sudeep Bhatia, "Variable time preference," *Cognitive psychology*, 2019, *111*, 53–79.

Heng, Joseph A, Michael Woodford, and Rafael Polania, "Efficient sampling and noisy decisions," *Elife*, 2020, *9*, e54962.

Hollingworth, Harry Levi, "The central tendency of judgment," *The Journal of Philosophy, Psychology and Scientific Methods*, 1910, *7* (17), 461–469.

Imas, Alex, Michael Kuhn, and Vera Mironova, "Waiting to choose: The Role of Deliberation in Intertemporal Choice," *American Economic Journal: Microeconomics*, 2021.

Jackson, Matthew O and Leeat Yariv, "Present bias and collective dynamic choice in the lab," *American Economic Review*, 2014, *104* (12), 4184–4204.

Kable, Joseph W and Paul W Glimcher, "An "as soon as possible" effect in human intertemporal decision making: behavioral evidence and neural mechanisms," *Journal of neurophysiology*, 2010, *103* (5), 2513–2531.

Khaw, Mel Win, Ziang Li, and Michael Woodford, "Cognitive imprecision and small-stakes risk aversion," *The review of economic studies*, 2021, *88* (4), 1979–2013.

Krajbich, Ian, Dingchao Lu, Colin Camerer, and Antonio Rangel, "The attentional drift-diffusion model extends to simple purchasing decisions," *Frontiers in psychology*, 2012, *3*, 193.

Laibson, David, "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443–478.

Loewenstein, George and Drazen Prelec, "Anomalies in intertemporal choice: Evidence and an interpretation," *The Quarterly Journal of Economics*, 1992, *107* (2), 573–597.

Lu, Jay and Kota Saito, "Random intertemporal choice," *Journal of Economic Theory*, 2018, *177*, 780–815.

Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa, "Failures in contingent reasoning: The role of uncertainty," *American Economic Review*, 2019, *109* (10), 3437–74.

Martino, Benedetto De, Sebastian Bobadilla-Suarez, Takao Nouguchi, Tali Sharot, and Bradley C Love, "Social information is integrated into value and confidence judgments according to its reliability," *Journal of Neuroscience*, 2017, *37* (25), 6066–6074.

_ , **Stephen M Fleming, Neil Garrett, and Raymond J Dolan**, "Confidence in value-based choice," *Nature neuroscience*, 2013, *16* (1), 105–110.

Mazur, James E, "An adjusting procedure for studying delayed reinforcement," *Commons, ML.; Mazur, JE.; Nevin, JA*, 1987, pp. 55–73.

Nielsen, Kirby and John Rehbeck, "When Choices are mistakes," *Available at SSRN 3481381*, 2020.

Oprea, Ryan, "What is complex," *American Economic Review*, 2020.

Polania, Rafael, Michael Woodford, and Christian C Ruff, "Efficient coding of subjective value," *Nature neuroscience*, 2019, *22* (1), 134–142.

Read, Daniel, "Is time-discounting hyperbolic or subadditive?," *Journal of risk and uncertainty*, 2001, *23* (1), 5–32.

_ , **Shane Frederick, Burcu Orsel, and Juwaria Rahman**, "Four score and seven years from now: The date/delay effect in temporal discounting," *Management Science*, 2005, *51* (9), 1326–1335.

Regenwetter, Michel, Daniel R Cavagnaro, Anna Popova, Ying Guo, Chris Zwilling, Shiau Hong Lim, and Jeffrey R Stevens, "Heterogeneity and parsimony in intertemporal choice.," *Decision*, 2018, *5* (2), 63.

Rubinstein, Ariel, ""Economics and psychology"? The case of hyperbolic discounting," *International Economic Review*, 2003, *44* (4), 1207–1216.

Simonson, Itamar and Amos Tversky, "Choice in context: Tradeoff contrast and extremeness aversion," *Journal of marketing research*, 1992, *29* (3), 281–295.

Soutschek, Alexander, Marius Moisa, Christian C Ruff, and Philippe N Tobler, "Frontopolar theta oscillations link metacognition with prospective decision making," *Nature Communications*, 2021, *12* (1), 1–8.

Sozou, Peter D, "On hyperbolic discounting and uncertain hazard rates," *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 1998, *265* (1409), 2015–2020.

Sprenger, Charles, "Judging experimental evidence on dynamic inconsistency," *American Economic Review*, 2015, *105* (5), 280–85.

Thaler, Richard, "Some empirical evidence on dynamic inconsistency," *Economics letters*, 1981, *8* (3), 201–207.

Tversky, Amos and Daniel Kahneman, "Judgment under uncertainty: Heuristics and biases," *science*, 1974, *185* (4157), 1124–1131.

Weitzman, Martin L, "Gamma discounting," *American Economic Review*, 2001, *91* (1), 260–271.

Woodford, Michael, "Modeling imprecision in perception, valuation, and choice," *Annual Review of Economics*, 2020, *12*, 579–601.

Xiang, Yang, Thomas Graeber, Benjamin Enke, and Samuel Gershman, "Confidence and Central Tendency in Perceptual Judgment," *Attention, Perception and Psychophysics*, 2021.

ONLINE APPENDIX

A Derivations for Bayesian Cognitive Noise Model

A.1 Model setup

Below we discuss the main behavioral predictions of a Bayesian cognitive imprecision model as outlined in Section 3. Suppose the DM has access to a mental simulation of the optimal action that we conceptualize as a "cognitive signal" S. We assume that S is an unbiased estimate of a^*, the normalized indifference point[30], and follows a scaled binomial distribution,

$$S \sim \frac{1}{n_2} Bin(n_2, a^*), \tag{8}$$

such that $0 < S < 1$. The parameter n_2 controls the precision of the mental simulation. The subjective likelihood of the utility-maximizing action based on a randomly drawn internal representation $\{S = s\}$ can then be represented by a binomial distribution:

$$\mathscr{L}(a^*|S = s) = P(S = s|a^*, n_2) = \binom{n_2}{sn_2} (a^*)^{sn_2}(1-a^*)^{(1-s)n_2}. \tag{9}$$

The DM holds a prior about his utility-maximizing action, A, which we broadly think of as the mathematical analogue of a decision maker's initial reaction to a choice problem. We assume that this prior can be represented by a Beta distribution, $A \sim Beta(n_1 d\ ,\ n_1(1-d))$. Here, d is the prior mean and carries the interpretation of a "cognitive default" action that the DM would take before deliberating about the problem. The parameter n_1, on the other hand, reflects the DM's confidence in (or precision of) their prior.[31] Note that the default action represents a *fraction* of the larger-later consumption, rather than an absolute consumption level.

A Bayesian DM accounts for the noisiness of his mental simulation by implicitly forming a posterior assessment of the utility-maximizing action. Given a Beta-distributed prior and a Binomial signal, this posterior belief, $A\Big|\{S = s\}$, is also Beta-distributed.[32] The DM's observed action given a mental signal is assumed to be the posterior mean:[33]

$$a^o = E[A\Big|\{S = s\}] = \lambda s + (1-\lambda)d \quad \text{with} \quad \lambda = n_2/(n_1 + n_2). \tag{10}$$

[30]This interpretation is possible since $u(c) = c^\alpha$.

[31]Note that $n_1 = a + b$ is a re-parameterization of the typical shape parameters a and b of the Beta distribution. n_1 is inversely related to the variance of the prior, $\sigma_A^2 = \frac{d\cdot(1-d)}{1+n_1}$.

[32]Specifically, $A_{S=s} \sim\ Beta(sn_2 + n_1 d\ ,\ n_2(1-s) + n_1(1-d))$.

[33]Focusing on the posterior mean is without much loss in the present context because the mean of a Beta(a,b) variable is $a/(a+b)$, the mode is $(a-1)/(a+b-2)$ and the median lies between the two.

This endogenizes the decision rule we posited in equation (2) of Section 3. Crucially, a more precise mental simulation (higher n_2) has a direct, negative effect on the weighting factor λ, which implies a lower weight on the cognitive default action. In the following subsection, we will thus focus on deriving behavioral predictions for changes in λ. In subsection A.3, we characterize cognitive uncertainty in the context of this model.

A.2 Derivations for behavioral predictions

All theorems and derivations in this subsection will solely concern a given subject's mean observed action, i.e., their average response aggregating across many unbiased signals. Given $\mathbb{E}[S] = a^*$, we define:

$$a^e := \mathbb{E}[a^o] = \lambda \cdot a^* + (1 - \lambda) \cdot d. \tag{11}$$

where $\lambda \in [0, 1]$ is our representation of cognitive precision. We derive the following behavioral predictions under the assumption of a canonical exponential discount function, $D(t) = \delta^t$, in order to disentangle the effect of shrinkage to a default from the predictions generated by present bias. Hence, $a^* = u^{-1}(\delta^{\Delta t})$ with $u(c) = c^\alpha, \alpha > 0$. Since $\alpha > 0$, we may allow δ to absorb α and in effect take $\alpha = 1$ in the proofs. As before (Section 6.2.2), we define:

$$RRR := \ln(\frac{1}{a^e}) = -\ln(a^e), \qquad \delta_H := e^{-\frac{RRR}{\Delta t}} \tag{12}$$

as the required rate of return and implied annualized discount factor. We will use the required rate of return per unit of time,

$$r := \frac{RRR}{\Delta t}, \tag{13}$$

as our measure of per-period impatience. We define short horizons as those time horizons where an exponential discounter behaves more patiently than a subject playing the default action:

$$SH := \{\Delta t \mid a^* > d\} \tag{14}$$

Long horizons, LH, are similarly defined by:

$$LH := \{\Delta t \mid a^* < d\} \tag{15}$$

We now turn to the theoretical predictions underlying the pre-registered predictions spelled out in Section 3.

Theorem 1 (Impatience over different time horizons)**.**

(i) Higher cognitive precision leads to less per-period impatience over short horizons.

$$\frac{\partial r}{\partial \lambda}\Big|_{\Delta t \in SH} < 0 \tag{16}$$

(ii) Higher cognitive precision leads to more per-period impatience over long horizons.

$$\frac{\partial r}{\partial \lambda}\Big|_{\Delta t \in LH} > 0 \tag{17}$$

Proof. Note that:

$$\frac{\partial a^e}{\partial \lambda} = a^* - d, \tag{18}$$

by definition. Hence, the sign of eq. (18) depends on whether it is evaluated over a short or long time horizon. We may now differentiate:

$$\frac{\partial r}{\partial \lambda} = \frac{1}{\Delta t}\frac{\partial RRR}{\partial \lambda} \tag{19}$$

$$= -\frac{1}{a^e \Delta t}\frac{\partial a^e}{\partial \lambda} \tag{20}$$

Since we trivially have $\Delta t, a^e > 0$, the sign of $\partial r/\partial \lambda$ is given by eq. (18) and the definitions (14) and (15), which yields the result. □

We note the trivial corollary that delivers Prediction 1 in the main text:

Corollary 1.1. *Subjects with perfect cognitive precision, $\lambda = 1$, have less pronounced short run impatience than those with imperfect cognitive precision, whereas the opposite is true concerning long run impatience.*

Given our measure of per-period impatience, we may show that per-period impatience decreases in the time delay (Δt).

Proposition 1 (Decreasing per-period impatience)**.**

(i) For those with complete cognitive precision, $\lambda = 1$, per-period impatience is constant in the time delay. Formally,

$$\frac{\partial r}{\partial \Delta t}\Big|_{\lambda=1} = 0 \tag{21}$$

(ii) For those with imperfect cognitive precision, $\lambda \in (0, 1)$, per-period impatience decreases in the time delay. In other words:

$$\frac{\partial r}{\partial \Delta t}\Big|_{\lambda<1} < 0 \tag{22}$$

Proof. We will consider the two cases: (i) $\lambda = 1$; (ii) $\lambda \in [0,1)$ separately.

In the case $\lambda = 1$ we trivially note that:

$$r = \frac{-\ln|a^e|}{\Delta t} = -\ln|\delta| \tag{23}$$

Hence, we have that:

$$\frac{\partial r}{\partial \Delta t} = 0 \tag{24}$$

For $\lambda \in (0,1)$, note that $\delta < 1$ implies:

$$\frac{\partial RRR}{\partial \Delta t} = -\frac{\lambda \delta^{\Delta t} \ln|\delta|}{a^e} > 0 \tag{25}$$

$$\frac{\partial^2 RRR}{(\partial \Delta t)^2} = -\frac{\lambda(1-\lambda)\delta^{\Delta t} d \ln^2|\delta|}{(a^e)^2} < 0 \tag{26}$$

meaning that the RRR is concave in Δt. The following expression describes how the RRR per unit of time changes in the time delay:

$$\frac{\partial r}{\partial \Delta t} = \frac{\frac{\partial RRR}{\partial \Delta t}\Delta t - RRR}{\Delta t^2} \tag{27}$$

A sufficient condition for (27) to have negative sign is therefore:

$$\Delta t \cdot \frac{\partial RRR}{\partial \Delta t} < RRR \tag{28}$$

We may now define the function:

$$g := RRR - \frac{\partial RRR}{\partial \Delta t}\Delta t \tag{29}$$

and differentiate to find:

$$\frac{\partial g}{\partial \Delta t} = -\frac{\partial^2 RRR}{(\partial \Delta t)^2}\Delta t \geq 0 \tag{30}$$

We note that at $\Delta t = 0$ we have:

$$g(0) = RRR(0) = -\ln|\lambda + (1-\lambda)d| > 0 \tag{31}$$

since $0 < d, \lambda < 1$. Hence, we find that g is positive for all $\Delta t > 0$:

$$g > 0 \; \Big| \Delta t > 0 \tag{32}$$

substituting in the definition of g shows that (28) is satisfied yielding the result. □

The following corollary underlies Prediction 2 in the main text:

Corollary 1.1. *The magnitude of per-period impatience's decrease in the time delay is smaller for those with perfect cognitive precision than for those with imperfect cognitive precision. Locally, this provides:*

$$\frac{\partial^2 r}{\partial \lambda \partial \Delta t}\Big|_{\lambda=1} > 0 \tag{33}$$

Proof. Note that the previous proposition provides that $dr/d\Delta t < 0$ for $\lambda < 1$ and is equal to zero for $\lambda = 1$. The result follows. □

It is important to note that the above theorems make no assumptions concerning the start time t_1 or end time t_2; but rather, only depend on the time delay $\Delta t = t_2 - t_1$. This is in line with our Predictions 1 and 2, which cover both delays starting in the present and in the future.

Next, we turn to the phenomenon of subadditivity. Subadditivity arises purely as a result of cognitive noise – as is well-known, $\beta - \delta$ preferences do not generate subadditivity.

Theorem 2 (Subadditivity). *Those subjects reporting cognitive uncertainty and an interior default will exhibit subadditivity in their choices. In other words, for $\lambda \in [0,1), d \in (0,1)$ we claim:*

$$SA := (RRR_{t_1,t_2} + RRR_{t_2,t_3}) - RRR_{t_1,t_3} > 0 \tag{34}$$

Since *RRR* only depends on the time delay and not the start time, it will be convenient to replace t_1, t_2, t_3 with the variables:

$$\Delta t_1 := t_2 - t_1, \qquad \Delta t_2 := t_3 - t_2$$

Taking a^e as a function of the time delay, our subadditivity condition can be rewritten as:

$$SA > 0 \tag{35}$$

$$\ln \left| \frac{a^e(\Delta t_1 + \Delta t_2)}{a^e(\Delta t_1) a^e(\Delta t_2)} \right| > 0 \tag{36}$$

$$\frac{a^e(\Delta t_1 + \Delta t_2)}{a^e(\Delta t_1) a^e(\Delta t_2)} > 1 \tag{37}$$

$$a^e(\Delta t_1 + \Delta t_2) > a^e(\Delta t_1) a^e(\Delta t_2) \tag{38}$$

Before we proceed to the proof, let us illustrate the effect of cognitive uncertainty on strict subadditivity through consideration of the edge cases $\lambda \in \{0,1\}, d \in \{0,1\}$. With perfect cognitive precision, $\lambda = 1$, the model reduces to standard exponential discounting and using (38) shows that there is no subadditivity:

$$\delta^{\Delta t_1 + \Delta t_2} = \delta^{\Delta t_1} \cdot \delta^{\Delta t_2} \tag{39}$$

In the presence of no cognitive precision, $\lambda = 0$, using (38) shows that there is subadditivity for any interior cognitive default $d \in (0, 1)$:

$$d > d^2 \tag{40}$$

Having discussed the corner cases we now proceed to the proof.

Proof. By (38) the existence of subadditivity is equivalent to:

$$a^e(\Delta t_1 + \Delta t_2) > a^e(\Delta t_1)a^e(\Delta t_2) \tag{41}$$

$$\lambda\delta^{\Delta t_1+\Delta t_2} + (1-\lambda)d > [\lambda\delta^{\Delta t_1} + (1-\lambda)d][\lambda\delta^{\Delta t_2} + (1-\lambda)d] \tag{42}$$

$$> \lambda^2\delta^{\Delta t_1+\Delta t_2} + (1-\lambda)^2 d^2 + \lambda(1-\lambda)(\delta^{\Delta t_1} + \delta^{\Delta t_2}) \tag{43}$$

Gathering our terms to the left hand side we find:

$$(\lambda - \lambda^2)\delta^{\Delta t_1+\Delta t_2} + (1-\lambda)d(1-(1-\lambda)d) - (\lambda - \lambda^2)d(\delta^{\Delta t_1} + \delta^{\Delta t_2}) > 0 \tag{44}$$

Since $\lambda \neq 1$ we may divide both sides by $(1-\lambda)$ to yield:

$$\lambda(\delta^{\Delta t_1+\Delta t_2} - d(\delta^{\Delta t_1} + \delta^{\Delta t_2})) + d(1-(1-\lambda)d) > 0 \tag{45}$$

We may define a function, $g(d)$ by:

$$g(d) := \lambda(\delta^{\Delta t_1+\Delta t_2} - d(\delta^{\Delta t_1} + \delta^{\Delta t_2})) + d(1-(1-\lambda)d) \tag{46}$$

so that subadditivity is equivalent to $g > 0 \,\big|\, d \in (0, 1)$. We now prove this claim. Note that g is quadratic in d with negative second derivative:

$$\frac{\partial^2 g}{\partial d^2} = -2(1-\lambda) < 0 \tag{47}$$

Accordingly, its unique minima on an interval will be found on the boundary points of the interval. In our case the boundary points are $d \in \{0, 1\}$. At $d = 0$ we find:

$$g(0) = \lambda\delta^{\Delta t_1+\Delta t_2} > 0 \tag{48}$$

At $d = 1$ we have:

$$g(1) = \lambda(1 + \delta^{\Delta t_1+\Delta t_2} - \delta^{\Delta t_1} - \delta^{\Delta t_2}) \tag{49}$$

If we view $g(1)$ as function of Δt_1, with, $h(\Delta t_1) := g(1)$, then we may note that:

$$h(0) = 0 \tag{50}$$

$$\frac{dh}{d\Delta t_1} = \lambda(\delta^{\Delta t_2} - 1)\delta^{\Delta t_1} \ln|\delta| \geq 0 \tag{51}$$

since $\delta \in (0,1)$, $\Delta t_1 \geq 0$, $\Delta t_2 \geq 0$. Consequently, we see that $g(1) \geq 0$ and may conclude that $g > 0$ for $d \in (0,1)$. $\square$

The following corollary delivers Prediction 3 in the main text.

Corollary 2.1. *The magnitude of subadditive behavior is greater for those with lower cognitive precision than for those who are certain ($\lambda = 1$).*

Proof. For those who are certain, we have that $SA = 0$, whereas for those that exhibit any uncertainty we have $SA > 0$. $\square$

Theorem 3. *There are no front-end delay effects.*

Proof. As mentioned earlier, *RRR* is a function of the time delay, Δt, not the individual start and end times. This precludes the existence of front-end delay effects. Formally, for any $l > 0$,

$$\Delta FE := RRR_{0,t_2} - RRR_{l,t_2+l} = \ln\left(\frac{\lambda u^{-1}(\delta^{t_2}) + (1-\lambda)d}{\lambda u^{-1}(\delta^{t_2}) + (1-\lambda)d}\right) = 0. \tag{52}$$

$\square$

The corresponding corollary underlying Prediction 4 in the main text is:

Corollary 3.1. *An increase in cognitive precision doesn't affect front-end delay effects.*

A.3 Derivations for Cognitive Uncertainty Measure

As laid out in Section 3, the DM subjectively perceives his optimal action as a *distribution* conditional on his noisy signal. This means: while the agent's loss function induces him to play $a^o = \mathbb{E}[A\big|\{S=s\}]$, the underlying perceived posterior distribution of the optimal action is Beta-distributed:

$$A\Big|\{S=s\} \sim \ Beta\left(\underbrace{sn_2 + n_1 d}_{\equiv a},\ \underbrace{n_2(1-s) + n_1(1-d)}_{\equiv b}\right) \tag{53}$$

where n_2 is the signal precision. Now, let us restate our definition of cognitive uncertainty,

$$p_{CU} := \mathbb{P}(|A\big|\{S=s\} - \mathbb{E}[A\big|\{S=s\}]| > c), \tag{54}$$

for fixed constant c. The objective of this subsection is to establish that increases in signal precision decrease cognitive uncertainty. Below, we develop two sets of results about this relationship. First, Corollary 2.1 provides a limit argument showing that any desired decrease in cognitive uncertainty can be achieved by an increase in signal precision. Second, to shed light on the case with low signal precision, Theorem 4 shows that cognitive uncertainty decreases with signal precision when using the closest Gaussian approximation of the Beta distribution.

To begin, we prove:

Proposition 2. *$\forall \kappa > 0, \forall \varepsilon > 0, \exists N \in \mathbb{N}$ such that $p_{CU} < \varepsilon$ for $n_2 > N$.*

Proof. By Chebyshev's inequality we see that for any positive number, κ:

$$p_{CU} < \frac{\mathrm{Var}(A\big|\{S = s\})}{\kappa^2} \tag{55}$$

and, since $A\big|\{S = s\} \sim Beta(n_2 s + n_1 d, n_2(1-s) + n_1(1-d))$ its variance is found to be:

$$\mathrm{Var}(A\big|\{S = s\}) = \frac{(n_2 s + n_1 d)(n_2(1-s) + n_1(1-d)),}{(n_1 + n_2)^2(n_1 + n_2 + 1)} = O(n_2^{-1}) \tag{56}$$

Accordingly, we find:

$$\lim_{n_2 \to \infty} p_{CU} = 0 \tag{57}$$

Which in turn yields the proposition. □

This proposition yields the following corollary:

Corollary 2.1. *Holding the signal value constant $\{S = s\}$ and given a base level of signal precision, n_2, there exists a constant Δn such that a desired decrease in cognitive uncertainty may be accomplished by increasing the signal precision by more than Δn.*
Formally, given a base signal precision, n_2, and a desired decrease in cognitive uncertainty, $\delta \in (0, p_{CU})$, there exists a quantity, $\Delta n \in \mathbb{N}$, such that:

$$n_2\prime > n_2 + \Delta n \rightarrow p_{CU} - p_{CU}\prime > \delta \tag{58}$$

with $n_2\prime$ and $p_{CU}\prime$ being the new signal precision and cognitive uncertainty respectively.

Proof. Given a signal precision n_2 and cognitive uncertainty, p_{CU}, we may apply the proposition to $\varepsilon = p_{CU} - \delta$. We then find that $\Delta n = N - n_2$. The result follows. □

In essence, this corollary formally states the intuition that any desired decrease in cognitive uncertainty may be accomplished through an increase in signal precision.

For a better approximation in cases of low signal precision, it is necessary to develop approximations of the Beta distribution. One such approximation follows from the Central Limit Theorem. We first prove a useful result:

Proposition 3. *Let $B_i \sim Beta(a_i, b_i)$, $n_i = a_i + b_i$ and $\forall i \in \mathbb{N}$, $\frac{a_i}{a_i+b_i} = \mu$, then*

$$B_i \xrightarrow{d} \mathcal{N}\left(\mu, \frac{\mu(1-\mu)}{n_i}\right)$$

as $a_i, b_i \to \infty$.

To prove this proposition we require the following lemma:

Lemma 3.1. *Let $Y_n \sim Gamma(na, 1)$ then $Y_n \xrightarrow{d} \mathcal{N}(na, a)$ as $n \to \infty$.*

Proof. Since the sum of Gamma variables follows a Gamma distribution,[34] we see that Y_n has the same distribution as:

$$\tilde{X} = \sum_{i=1}^{n} X_i$$

where X_i are i.i.d. random variables sampled from $Gamma(a, 1)$.

Now, by the Central Limit Theorem, we have:

$$\sqrt{\frac{n}{a}}\left(\frac{\tilde{X}}{n} - a\right) \xrightarrow{d} \mathcal{N}(0,1) \tag{59}$$

which yields:

$$Y_n \xrightarrow{d} \mathcal{N}(na, na) \tag{60}$$

□

We are now in a position to prove our proposition.

Proof. Let $X_i \sim Gamma(a_i, 1)$ and $Y_i \sim Gamma(b_i, 1)$ be independent random variables. Then we know[35] that:

$$Z_i = g(X_i, Y_i) := \frac{X_i}{X_i + Y_i} \sim Beta(a_i, b_i)$$

Note, that if we scale both X_i and Y_i by $(a_i + b_i)^{-1}$ that Z_i remains unchanged. Further-

[34] Let $X_n \sim Gamma(a_n, 1)$ and $Y_n \sim Gamma(b_n, 1)$ be independent Gamma variables. Then $Z_n = X_n + Y_n \sim Gamma(a_n + b_n, 1)$.

[35] This may be verified by consideration of the joint density of X_i, Y_i; making the transformation $V = X_i + Y_i$, with W_i as defined earlier; and finding the marginal density of W_i.

more, by our lemma above, we have:

$$\begin{pmatrix} \frac{X_i}{a_i+b_i} \\ \frac{Y_i}{a_i+b_i} \end{pmatrix} \xrightarrow{d} \mathcal{N}\left(\begin{pmatrix} \frac{a_i}{a_i+b_i} \\ \frac{b_i}{a_i+b_i} \end{pmatrix}, \begin{pmatrix} \frac{a_i}{(a_i+b_i)^2} & 0 \\ 0 & \frac{b_i}{(a_i+b_i)^2} \end{pmatrix}\right) \tag{61}$$

$$\xrightarrow{d} \mathcal{N}\left(\begin{pmatrix} \mu \\ 1-\mu \end{pmatrix}, \begin{pmatrix} \frac{\mu}{n_i} & 0 \\ 0 & \frac{1-\mu}{n_i} \end{pmatrix}\right) \tag{62}$$

We now employ the Delta method. A Taylor expansion of $g(x,y)$ yields that to first order:

$$g(x,y) \approx g(x_0,y_0) + \nabla g(x_0, y+0)\cdot(x-x_0, y-y_0) \tag{63}$$

$$\approx g(x_0,y_0) + (\frac{y_0}{(x_0+y_0)^2}, \frac{-x_0}{(x_0+y_0)^2})\cdot(x-x_0,y-y_0) \tag{64}$$

Accordingly, we find that:

$$Z_i = g(X_i, Y_i) \tag{65}$$

$$\approx g(\mu, 1-\mu) + ((1-\mu), -\mu)\cdot(X_i-\mu, Y_i-(1-\mu)) \tag{66}$$

$$\xrightarrow{d} \mu + ((1-\mu),-\mu)\cdot\mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \frac{\mu}{n} & 0 \\ 0 & \frac{1-\mu}{n} \end{pmatrix}\right) \tag{67}$$

$$\xrightarrow{d} \mathcal{N}\left(\mu, \begin{pmatrix} 1-\mu & -\mu \end{pmatrix}\begin{pmatrix} \frac{\mu}{n} & 0 \\ 0 & \frac{1-\mu}{n} \end{pmatrix}\begin{pmatrix} 1-\mu \\ -\mu \end{pmatrix}\right) \tag{68}$$

$$\xrightarrow{d} \mathcal{N}\left(\mu, \frac{\mu(1-\mu)}{n_i}\right) \tag{69}$$

□

This proposition provides the simplest Gaussian approximation of the Beta distribution; however, this approximation may be put on more intuitive grounds by taking into account geometric aspects of density function. In this case, we will consider the peak and the points of inflection. For a Gaussian, the mode is found at μ and the points of inflection are found at $\mu \pm \sigma$. For the Beta distribution we have that the mode is given by:

$$m = \frac{a-1}{a+b-2} = \frac{n\mu-1}{n-2} \tag{70}$$

When $a, b > 2$ the Beta's density function is bell shaped and we may find the points of inflection about the mode. If we define the constant:

$$\kappa := \frac{1}{a+b-2}\sqrt{\frac{(a-1)(b-1)}{a+b-3}} \tag{71}$$

the points of inflection may be written as $m \pm \kappa$. Note that $\kappa/\sigma \to 1$ and $m \to \mu$ as $n \to \infty$. Accordingly, we may employ the following approximation:

$$Beta(a,b) \approx \mathcal{N}(m,\kappa) \tag{72}$$

The goodness of fit for this approximation for our purposes may shown empirically. Using this approximation, we may return to our original goal of demonstrating that cognitive uncertainty decreases in the signal precision and claim:

Proposition 4. *Let $N_1(m_1,\kappa_1^2)$ be the normal approximation to $A_1 \sim Beta(a_1,b_1)$ and $N_2(m_2,\kappa_2^2)$ the normal approximation to $A_2 \sim Beta(a_2,b_2)$ with: $a_1,a_2,b_1,b_2 > 2$, $\frac{a_1}{a_1+b_1} = \frac{a_2}{a_2+b_2}$ and $a_2 > a_1, b_2 > a_1$ then for fixed $c \in (0, min_{i=1,2}\{m_i, 1-m_i\})$:*

$$\mathbb{P}(|N_2 - m_2| < c) > \mathbb{P}(|N_1 - m_1| < c) \tag{73}$$

where m_i is the mode of A_i and κ_i is the distance from the mode to the points of inflection for A_i.

Proof. We note that:

$$\mathbb{P}(|N_i - m_i| < c) = \mathbb{P}(|Z| < c/\kappa_i) \tag{74}$$

and that κ as defined in (71) satisfies[36] $\partial\kappa/\partial a, \partial\kappa/\partial b < 0$ for $a,b > 2$. Accordingly, we have that $\kappa_2 < \kappa_1$. Hence, we see that:

$$\mathbb{P}(|N_1 - m_1| < c) = \mathbb{P}(|Z| < c/\kappa_1) < \mathbb{P}(|Z| < c/\kappa_2) = \mathbb{P}(|N_2 - m_2| < c) \tag{75}$$

□

This provides us with our final result:

Theorem 4. *Holding the signal $\{S = s\}$ constant, cognitive uncertainty decreases with increases the signal precision in the Gaussian approximation eq.* (72).

$$\frac{\Delta p_{CU}}{\Delta n_2} < 0 \tag{76}$$

Proof. Apply the previous proposition with respect to the values of a, b from eq. (53). □

[36]Just reparametrize it under $a = x+1, b = y+1$ and square the expression.

B Additional Figures

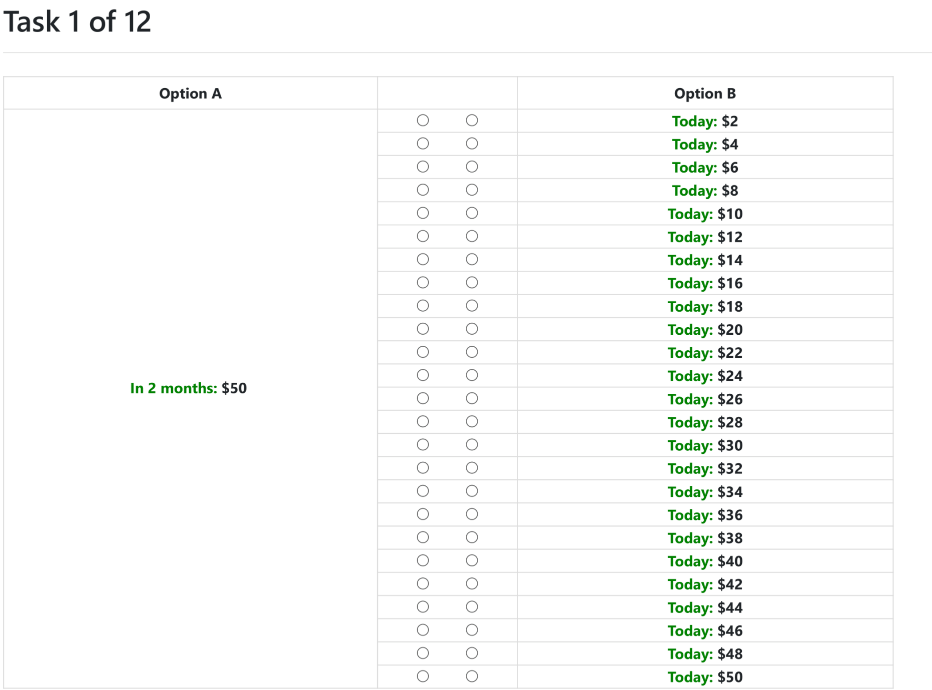

Figure 10: Screenshot of an example decision screen in *Money Main*

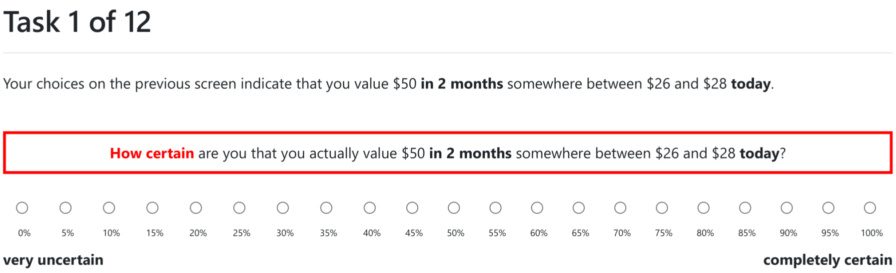

Figure 11: Screenshot of an example cognitive uncertainty elicitation screen in *Money Main*

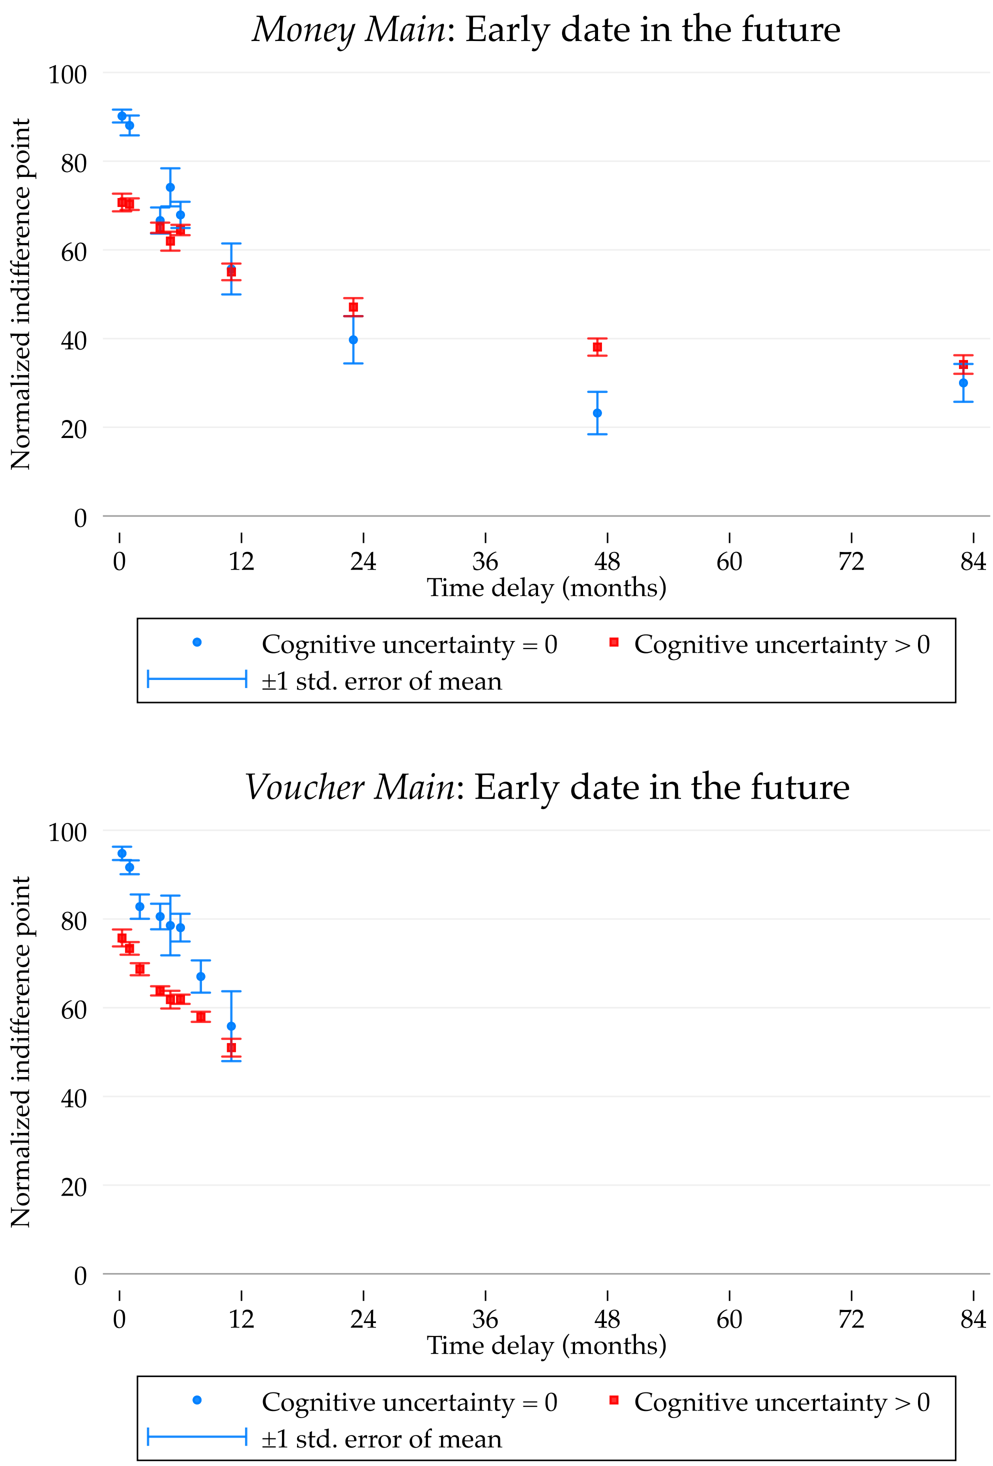

Figure 12: Observed discounting with $t_1 > 0$ in *Money Main* (top panel, $N = 2792$) and *Voucher Main*, $N = 2154$ (bottom panel). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

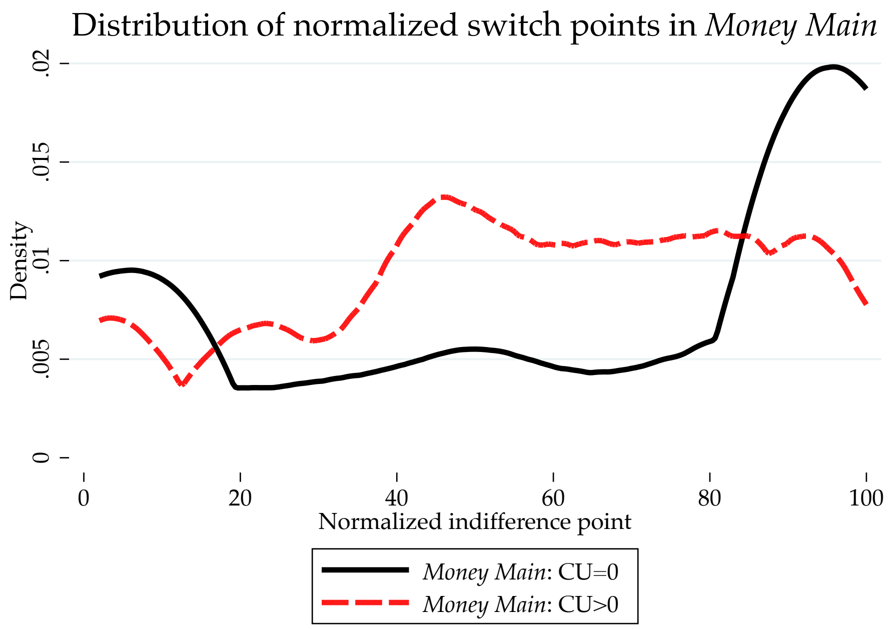

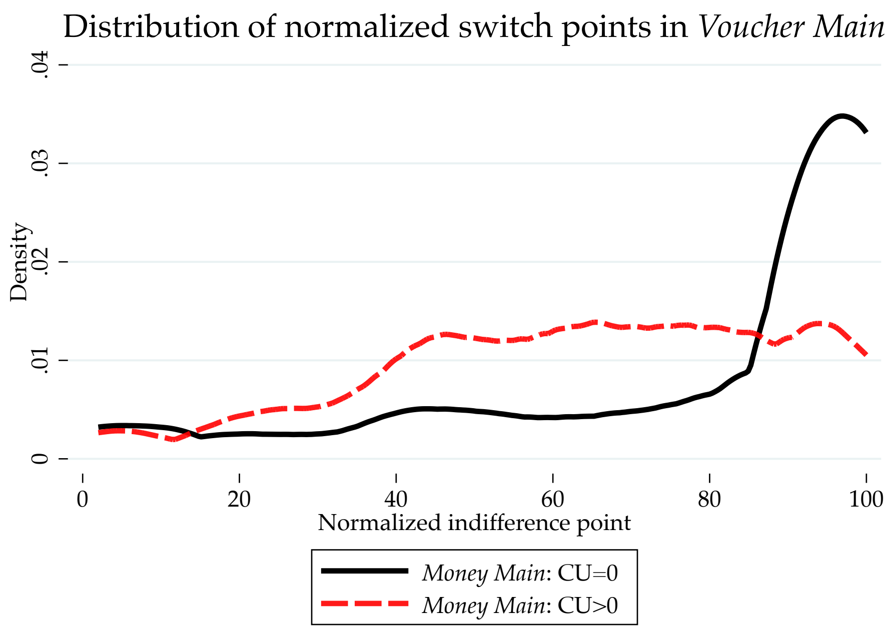

Figure 13: Kernel density plots of the distribution of normalized indifference points in the baseline experiments, separately for decisions that reflect zero or strictly positive cognitive uncertainty. Kernel is Epanechnikov.

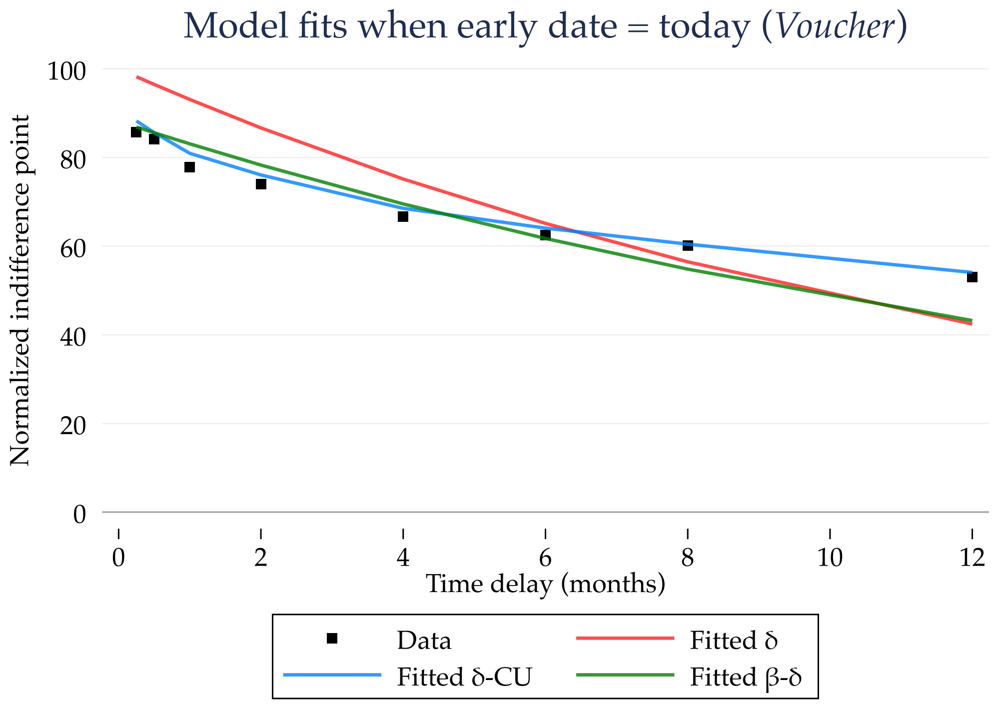

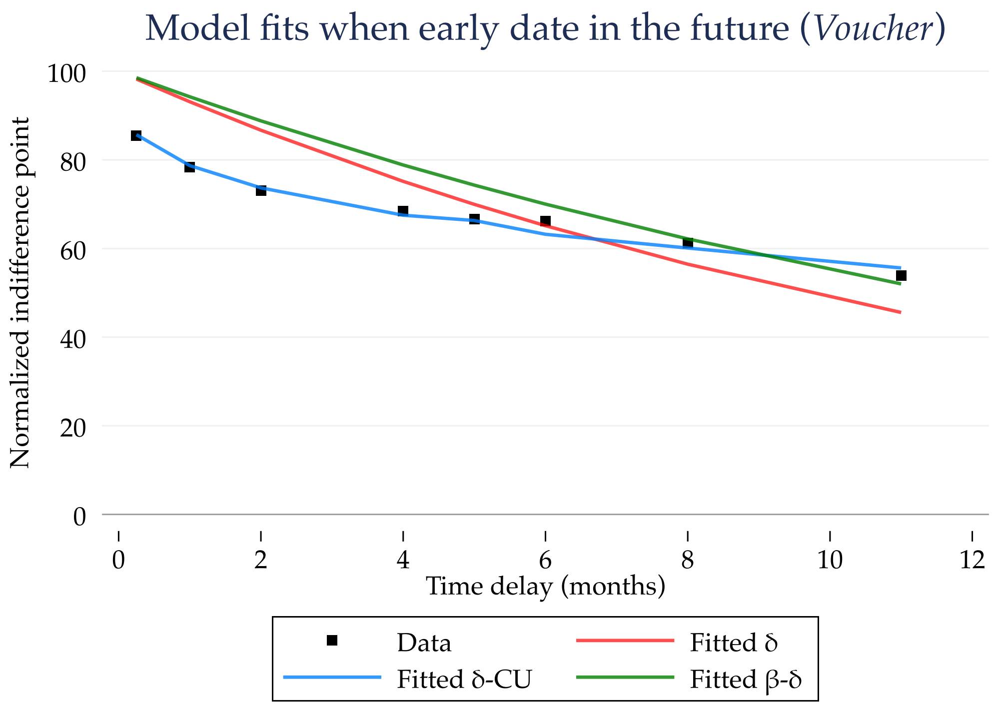

Figure 14: Model fit vs. data in *Voucher Main*. The model predictions are computed as fitted values of the parameter estimates in Table 4.

C Additional Tables

Table 5: List of treatment conditions

Treatment	**Description**	**Sample**	*Pre-registration*	**Covered in**
Money Main	Hypothetical money-early-versus-later experiments	645	`https://aspredicted.org/kg7zs.pdf`	Sections 4–6
Voucher Main	Incentivized UberEats Voucher Experiments	500	`https://aspredicted.org/b4pw2.pdf`	Sections 4–6
Money Complex Dates	Like *Voucher Main*, except with later payment date displayed as math expression	149	`https://aspredicted.org/77xp6.pdf`	Section 8.1
Money Complex Amounts	Like *Voucher Main*, except with payoff amount displayed as math expression	153	`https://aspredicted.org/77xp6.pdf`	Section 8.1
Money Load	Like *Voucher Main*, except with cognitive load manipulation	154	`https://aspredicted.org/77xp6.pdf`	Section 8.1
Money Main Replication	Like *Voucher Main*, used as control group for *Money Complex Dates*, *Money Complex Amounts* and *Money Load* (within-session randomization)	161	`https://aspredicted.org/77xp6.pdf`	Section 8.1
Voucher Advice	Like *Voucher Main*, except introducing piece of advice	153	`https://aspredicted.org/jk5s5.pdf`	Section 8.2
Money Within-Subject Load	Like *Money Main*, except across-round variation in cognitive load	400	`https://aspredicted.org/av2y2.pdf`	Appendix F

Notes. List of all treatments included in this paper.

Table 6: Cognitive uncertainty and across-trial choice variability

	Dependent variable: Abs. diff. b/w normalized indifference points					
Treatment:	*Money Main*			*Voucher Main*		
	(1)	(2)	(3)	(4)	(5)	(6)
Ave. cognitive uncertainty	0.14***	0.14***	0.14***	0.14***	0.13***	0.13***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Time delay FE	No	Yes	Yes	No	Yes	Yes
Demographic controls	No	No	Yes	No	No	Yes
Observations	1290	1290	1290	1000	1000	1000
R^2	0.03	0.04	0.04	0.03	0.04	0.04

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is computed as absolute difference between the normalized indifference points in two repetitions of the exact same choice list. The independent variable is average cognitive uncertainty across the two repetitions of the choice list. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Cognitive uncertainty and impatience over one week

	Dependent variable: Normalized indifference point							
Treatment:	*Money Main*				*Voucher Main*			
Sample:	$t1 = 0$		$t1 > 0$		$t1 = 0$		$t1 > 0$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cognitive uncertainty	-0.66***	-0.65***	-0.58***	-0.55***	-0.66***	-0.65***	-0.61***	-0.64***
	(0.10)	(0.11)	(0.11)	(0.10)	(0.13)	(0.13)	(0.16)	(0.14)
Payment amount FE	No	Yes	No	Yes	No	Yes	No	Yes
Round FE	No	Yes	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	350	350	218	218	404	404	152	152
R^2	0.20	0.23	0.20	0.30	0.15	0.18	0.21	0.34

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes decisions in which the time delay is given by one week. Columns (1)–(2) and (5)–(6) include those trials in which the early payment date is today, and columns (3)–(4) and (7)–(8) those in which the early payment date is in the future. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Cognitive uncertainty and increasing per-period patience

	Dependent variable: Implied per-period patience δ_H							
Treatment:	*Money Main*				*Voucher Main*			
Sample:	$t1 = 0$		$t1 > 0$		$t1 = 0$		$t1 > 0$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Time delay (years)	0.058***	0.058***	0.049***	0.050***	0.18***	0.18***	0.20***	0.20***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.03)	(0.03)	(0.05)	(0.05)
Time delay × Cognitive uncertainty	0.00076***	0.00074***	0.00059***	0.00056***	0.0059***	0.0056***	0.0062***	0.0063***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Cognitive uncertainty	-0.0028***	-0.0027***	-0.0027***	-0.0026***	-0.0064***	-0.0063***	-0.0070***	-0.0071***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Payment amount FE	No	Yes	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	4948	4948	2792	2792	3846	3846	2154	2154
R^2	0.20	0.21	0.16	0.18	0.11	0.13	0.12	0.12

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. Columns (1)–(2) and (5)–(6) include those trials in which the early payment date is today, and columns (3)–(4) and (7)–(8) those in which the early payment date is in the future. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Cognitive uncertainty and insensitivity to time delays: Including participant fixed effects

	Dependent variable: Normalized indifference point							
Treatment:	*Money Main*				*Voucher Main*			
Sample:	$t1 = 0$		$t1 > 0$		$t1 = 0$		$t1 > 0$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Time delay (years)	-6.99***	-6.96***	-6.77***	-6.81***	-33.9***	-34.0***	-30.7***	-30.9***
	(0.36)	(0.36)	(0.40)	(0.39)	(1.99)	(1.96)	(3.88)	(3.89)
Time delay × Cognitive uncertainty	0.055***	0.055***	0.044***	0.045***	0.25***	0.25***	0.30*	0.30*
	(0.01)	(0.01)	(0.01)	(0.01)	(0.07)	(0.07)	(0.16)	(0.16)
Cognitive uncertainty	-0.26***	-0.26***	-0.28***	-0.28***	-0.30***	-0.30***	-0.42***	-0.41***
	(0.04)	(0.04)	(0.05)	(0.05)	(0.06)	(0.06)	(0.09)	(0.08)
Payment amount FE	No	Yes	No	Yes	No	Yes	No	Yes
Round FE	No	Yes	No	Yes	No	Yes	No	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4948	4948	2792	2792	3846	3846	2154	2154
R^2	0.66	0.66	0.68	0.69	0.73	0.74	0.71	0.71

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. Columns (1)–(4) include data from *Money Main*, where columns (1)–(2) restrict attention to decision problems with $t_1 = 0$ and columns (3)–(4) to problems with $t_1 > 0$. An analogous logic applies to columns (5)–(8) for *Voucher Main*. Demographic controls include age, gender and income bucket. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Distribution of participant-level estimates of model parameters

	Money Main (MPL & Direct Elicitation)			
	(1)	(2)	(3)	(4)
	δ	$\beta-\delta$	$\delta-CU$	$\beta-\delta$ $-CU$
	Median (25 / 75 pctl.)	Median (25 / 75 pctl.)	Median (25 / 75 pctl.)	Median (25 / 75 pctl.)
$\hat{\delta}$	0.96 (0.90 / 0.99)	0.97 (0.92 / 0.99)	0.97 (0.91 / 0.99)	0.97 (0.92 / 0.99)
$\hat{\beta}$		0.88 (0.66 / 0.99)		0.96 (0.74 / 1.00)
$\hat{d}$			0.51 (0.27 / 0.73)	0.51 (0.25 / 0.73)

Notes. Distribution of estimates of different versions of eq. (7) estimated at the subject level. MPL = multiple price list. Each column corresponds to a separate model specification. Column (1): set $\beta = 1$ and $p_{CU} = 0$. Column (2): set $p_{CU} = 0$. Column (3): set $\beta = 1$. All estimations accommodate utility curvature: a representative-agent CRRA parameter of $\hat{\gamma} = 0.94$ was separately estimated on the risky choice data and used in the participant-level estimations on the intertemporal choice data. The exponential parameter δ is the monthly discount factor.

Table 11: Correlations between participant-level estimates of model parameters in β—δ—CU specification

	$\hat{\delta}$	$\hat{\beta}$	$\hat{d}$	Mean stated CU
$\hat{\delta}$	1.000			
$\hat{\beta}$	0.198*** (0.000)	1.000		
$\hat{d}$	-0.169*** (0.000)	-0.015 (0.705)	1.000	
Mean stated CU	-0.100* (0.011)	-0.024 (0.546)	-0.096* (0.015)	1.000

Notes. Pairwise correlations of participant-level estimates of equation (7). p-values shown in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Direct Elicitation Experiments

As part of our *Money Main* experiments, each subject completed six additional intertemporal choice problems that were administered in a direct elicitation format rather than using MPLs. That is, in each of these decisions, subjects were directly asked which monetary amount to be received in $t = t_1$ is worth as much to them as receiving $\$y_2$ in

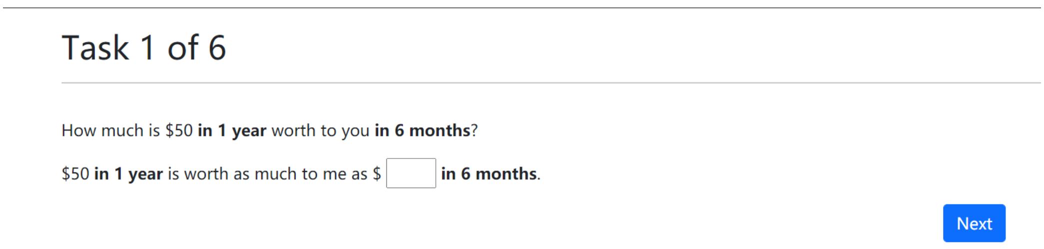

Figure 15: Screenshot of an example decision screen in the direct elicitation part of *Money Main*

$t = t_2$, see Figure 15 for an example screenshot.[37] After participants had indicated their indifference amount, the next screen again elicited cognitive uncertainty, see Figure 16.

We here replicate all of our main analyses using these direct elicitation data. First, Table 12 shows that cognitive uncertainty is again significantly correlated with the magnitude of across-trial inconsistencies (choice variability), as defined by the absolute difference in normalized indifference points across two repetitions of the same question.

Second, Table 13 documents that cognitive uncertainty is strongly and significantly correlated with impatience over a horizon of one week. Third, columns (1)–(2) of Table 14 document that cognitive uncertainty is highly predictive of a reduced sensitivity of intertemporal choice behavior with respect to variation in the time delay, as we can infer from the significant interaction term. Columns (3)–(4) show the same patterns by documenting that cognitive uncertainty is strongly predictive of decreasing impatience as the time delay increases, as we can again infer from the significant interaction term. Figure 17 visualizes these patterns.

Next, Table 15 documents that subadditivity effects strongly increase in cognitive uncertainty, see columns (2)–(3), (5)–(6) and (8)–(9). Indeed, as we can see from the usually insignificant raw term "1 if long interval", there is no significant evidence for subadditivity among subjects who indicate cognitive uncertainty of zero.

Finally, Table 16 replicates the result that cognitive uncertainty is uncorrelated with front-end delay effects. This again highlights that "not anything goes" but that cognitive uncertainty is only predictive of a specific set of empirical regularities as pre-registered.

[37]The only difference between the choice problems in the direct elicitation experiments and the MPL is that (to save time) we only elicited direct elicitation problems in which the early payment date was today.

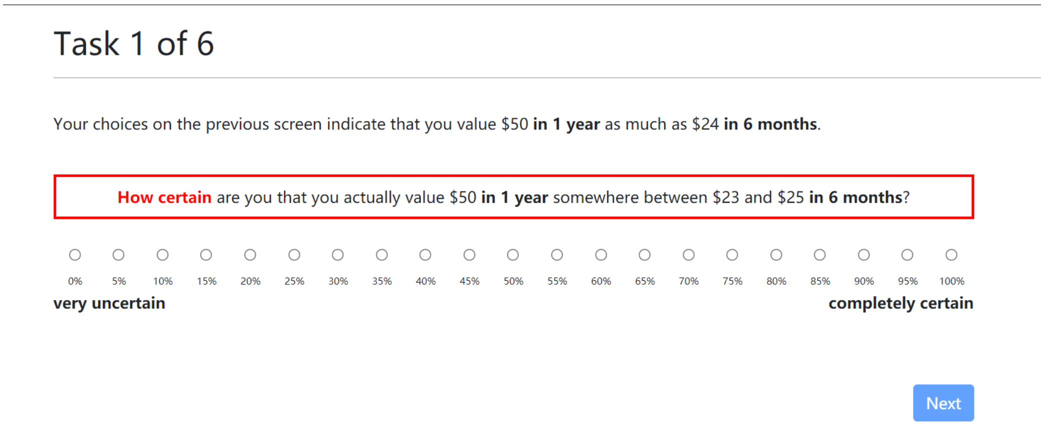

Figure 16: Screenshot of an example cognitive uncertainty elicitation screen in the direct elicitation part of *Money Main*

Table 12: Cognitive uncertainty and across-trial choice variability: Direct elicitation

	Dependent variable:		
	Abs. diff. b/w normalized indiff. points		
	(1)	(2)	(3)
Ave. cognitive uncertainty	0.083**	0.079**	0.079**
	(0.03)	(0.04)	(0.04)
Time delay FE	No	Yes	Yes
Demographic controls	No	No	Yes
Observations	645	645	645
R^2	0.01	0.02	0.02

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The dependent variable is computed as absolute difference between the normalized indifference points in two repetitions of the exact same choice list. The independent variable is average cognitive uncertainty across the two repetitions of the choice task. All observations are from the direct elicitation experiments. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 13: Cognitive uncertainty and impatience over one week: Direct elicitation

	Dependent variable: Normalized indifference point		
	(1)	(2)	(3)
Cognitive uncertainty	-0.59***	-0.59***	-0.57***
	(0.10)	(0.11)	(0.11)
Payment amount FE	No	Yes	Yes
Demographic controls	No	No	Yes
Observations	327	327	327
R^2	0.13	0.17	0.17

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes decisions in which the time delay is given by one week. All observations are from the direct elicitation experiments. In these experiments, the early payment date is always today. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: Cognitive uncertainty and diminishing impatience: Direct elicitation

	Dependent variable: Normalized indifference point		Implied per-period patience δ_H	
	(1)	(2)	(3)	(4)
Time delay (years)	-7.16***	-7.08***	0.043***	0.044***
	(0.49)	(0.48)	(0.00)	(0.00)
Time delay × Cognitive uncertainty	0.091***	0.084***	0.0011***	0.0011***
	(0.02)	(0.01)	(0.00)	(0.00)
Cognitive uncertainty	-0.43***	-0.40***	-0.0050***	-0.0047***
	(0.05)	(0.05)	(0.00)	(0.00)
Payment amount FE	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes
Observations	3870	3870	3870	3870
R^2	0.17	0.19	0.17	0.19

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. All observations are from the direct elicitation experiments. In these experiments, the early payment date is always today. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

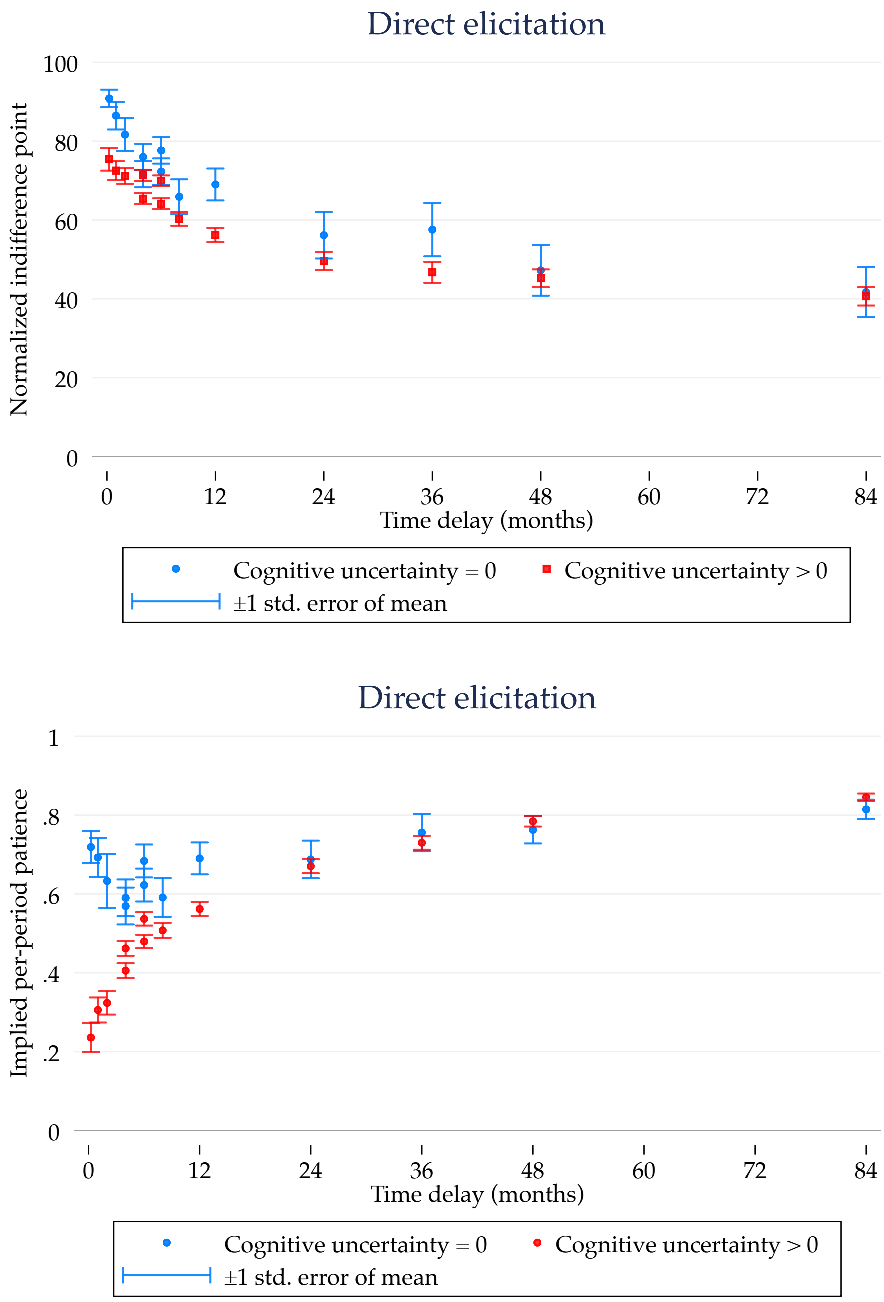

Figure 17: Observed discounting in the direct elicitation experiments (top panel, $N = 4,614$). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). Per-period patience is computed as $\delta_H(a^o) \equiv e^{-RRR/\Delta t} = (a^o)^{1/\Delta t}$, where a^o is the observed normalized indifference point.

Table 15: Cognitive uncertainty and subadditivity: Direct elicitation

	Dependent variable: Composite indifference point								
Sample:	Full			Set 1			Set 2		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1 if one long interval	7.45***	3.81***	3.93***	7.05***	3.72*	3.72*	7.86***	3.66*	3.77*
	(0.70)	(1.40)	(1.41)	(1.05)	(1.94)	(1.94)	(0.93)	(2.02)	(2.04)
1 if one long interval × Cognitive uncertainty		0.20***	0.19***		0.18**	0.18**		0.23***	0.22**
		(0.06)	(0.06)		(0.08)	(0.08)		(0.09)	(0.09)
Cognitive uncertainty		-0.47***	-0.47***		-0.53***	-0.53***		-0.41***	-0.39***
		(0.07)	(0.07)		(0.10)	(0.10)		(0.10)	(0.11)
Set FE	Yes	Yes	Yes	No	No	No	No	No	No
Payment amount FE	No	No	Yes	No	No	Yes	No	No	Yes
Observations	1290	1290	1290	654	654	654	636	636	636
R^2	0.02	0.06	0.07	0.01	0.08	0.10	0.02	0.05	0.07

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. All observations are from the direct elicitation experiments. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 16: Cognitive uncertainty and front-end delay effects: Direct elicitation

	Dependent variable:			
	Normalized indifference point			
	(1)	(2)	(3)	(4)
1 if front end delay	5.54***	5.28***	5.26***	5.27***
	(0.69)	(1.32)	(1.32)	(1.31)
Front-end delay × Cognitive uncertainty		0.060	0.061	0.056
		(0.06)	(0.06)	(0.06)
Cognitive uncertainty		-0.35***	-0.35***	-0.32***
		(0.06)	(0.06)	(0.06)
Set FE	Yes	Yes	Yes	Yes
Payment amount FE	No	No	Yes	Yes
Demographic controls	No	No	No	Yes
Observations	1290	1290	1290	1290
R^2	0.01	0.06	0.07	0.10

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. All observations are from the direct elicitation experiments. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Complexity and Load Experiments

E.1 Screenshots of Decision Screens

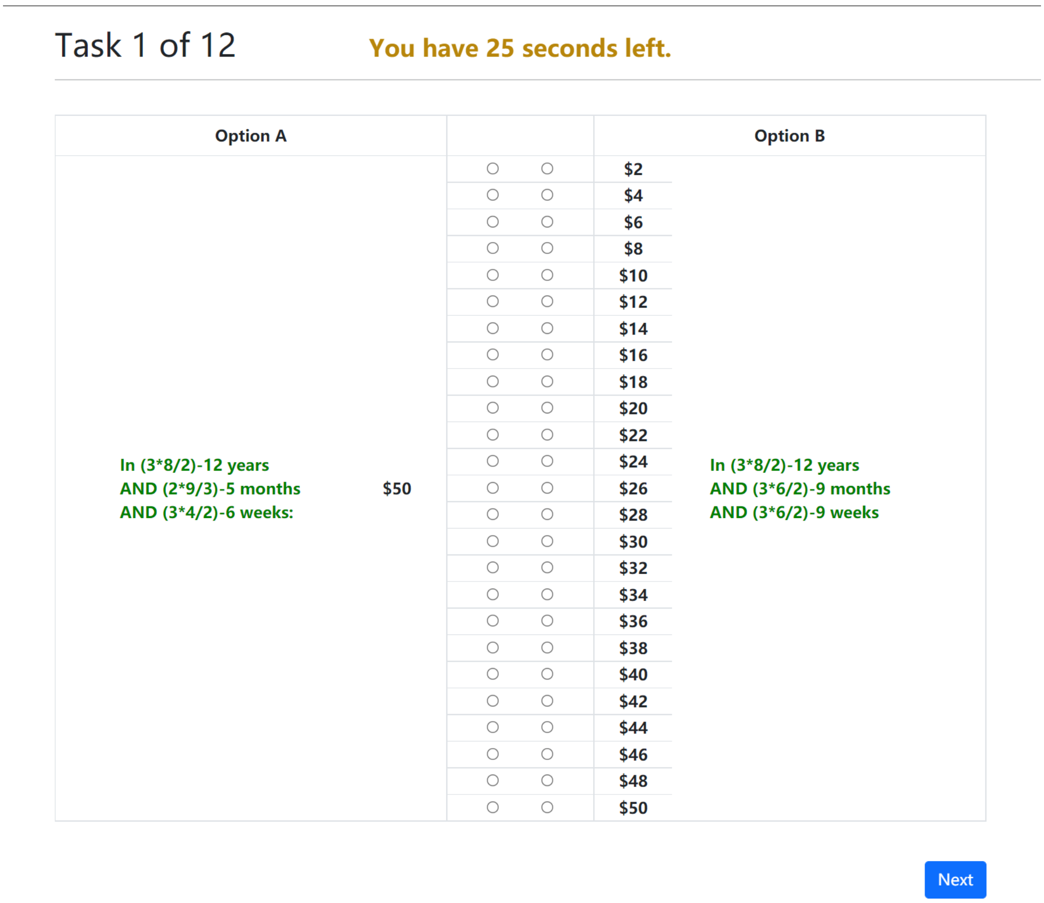

Figure 18: Screenshot of an example decision screen in *Money Complex Dates*

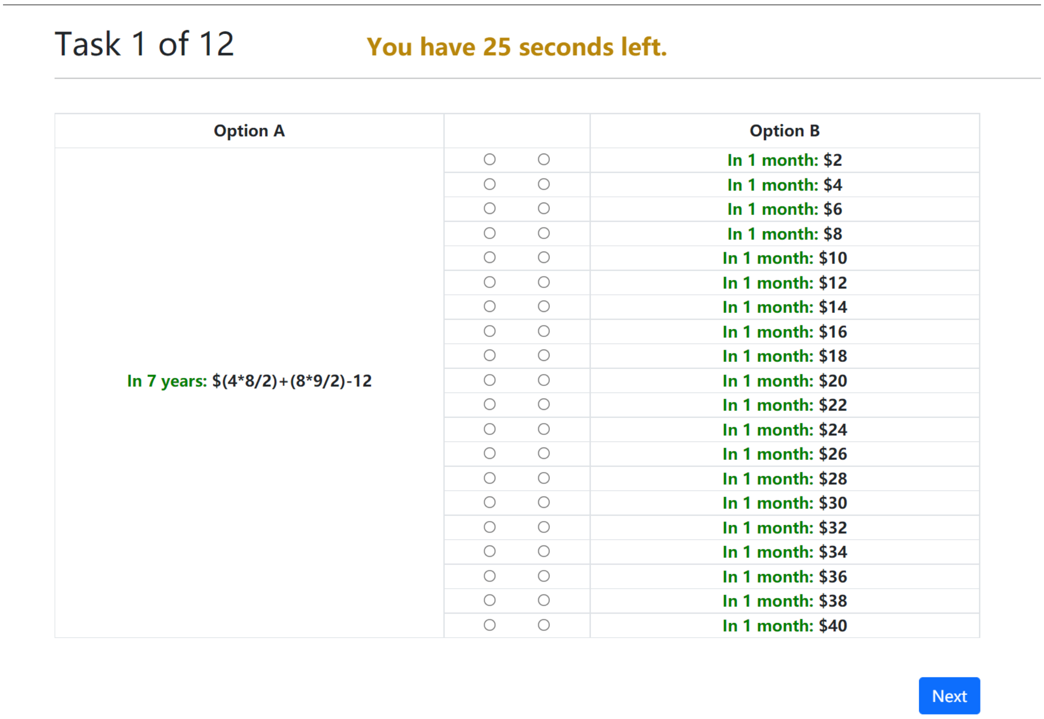

Figure 19: Screenshot of an example decision screen in *Money Complex Amounts*

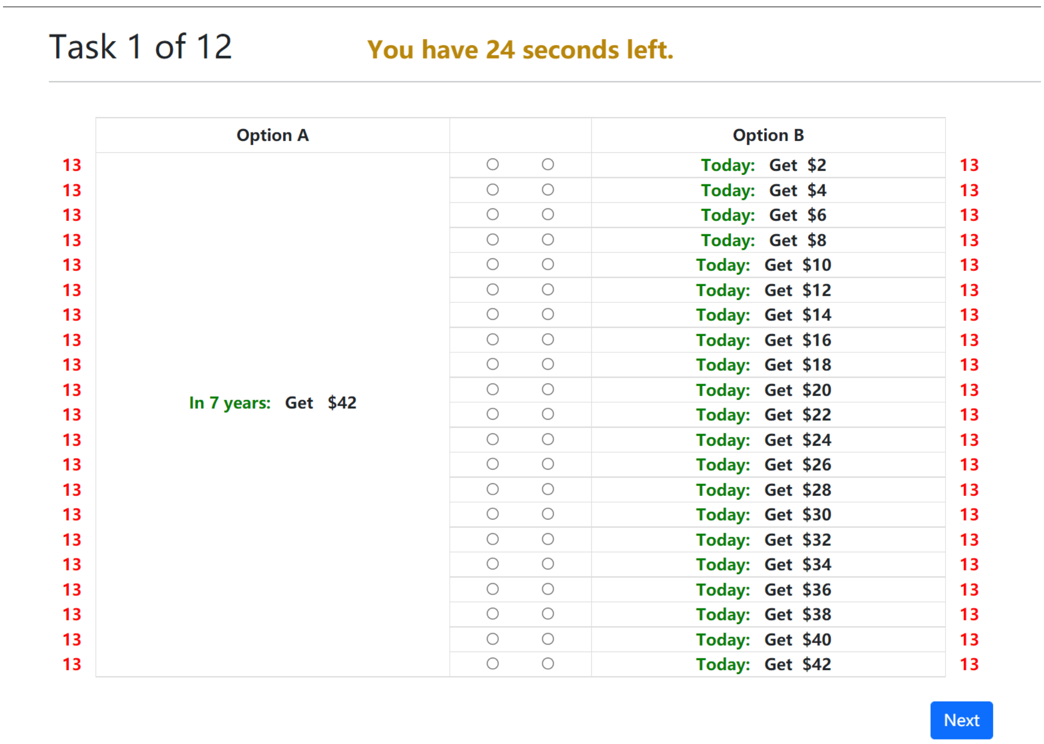

Figure 20: Screenshot of an example decision screen in *Money Load*

E.2 Results

Table 17 summarizes the results for all treatments that manipulate either complexity or cognitive load. For each of the three treatment variations, we compare behavior to the treatment *Money Main replication*, which was administered along with the complexity and load treatments. Columns (1)–(3) show the results for choice problems in which the early date is today, while columns (4)–(6) summarize analogous results for $t_1 > 0$. Througout, the positive interaction coefficients between the time delay and the more complex / load treatments indicate that people's decisions are more inelastic to the time delay when the choice problems are more complex or they are placed under cognitive load. Figures 21–25 visualize these patterns.

Table 17: Complexity and load manipulations

	Dependent variable: Normalized indifference point					
Sample:	$t1 = 0$			$t1 > 0$		
	(1)	(2)	(3)	(4)	(5)	(6)
Time delay (years)	-4.97***	-4.94***	-4.89***	-4.84***	-4.84***	-4.85***
	(0.55)	(0.54)	(0.55)	(0.62)	(0.62)	(0.62)
1 if *Complex Dates*	3.17			1.38		
	(3.00)			(2.94)		
Time delay × 1 if *Complex Dates*	2.97***			3.36***		
	(0.79)			(0.88)		
1 if *Complex Amounts*		0.86			-2.63	
		(2.91)			(3.00)	
Time delay × 1 if *Complex Amounts*		2.07***			2.43***	
		(0.75)			(0.84)	
1 if *Load*			-2.34			-2.68
			(3.04)			(3.04)
Time delay × 1 if *Load*			1.64**			2.00**
			(0.78)			(0.82)
Payment amount FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2381	2405	2428	1339	1363	1352
R^2	0.08	0.08	0.10	0.07	0.06	0.07

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. Columns (1)–(6) include data from *Money Main Replication*, *Money Complex Dates*, *Money Complex Amounts* and *Money Load*. Columns (1)–(3) restrict attention to decision problems with $t_1 = 0$ and columns (4)–(6) to problems with $t_1 > 0$. Demographic controls include age, gender and income bucket. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

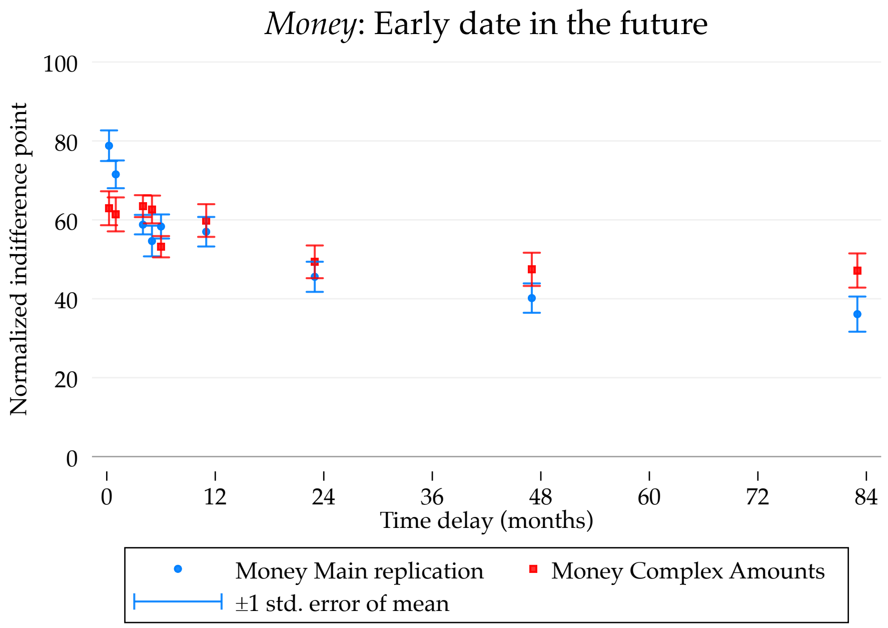

Figure 21: Observed discounting with $t_1 > 0$ in *Money Main replication* ($N = 161$ subjects) and *Money Complex Amounts* ($N = 153$ subjects). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

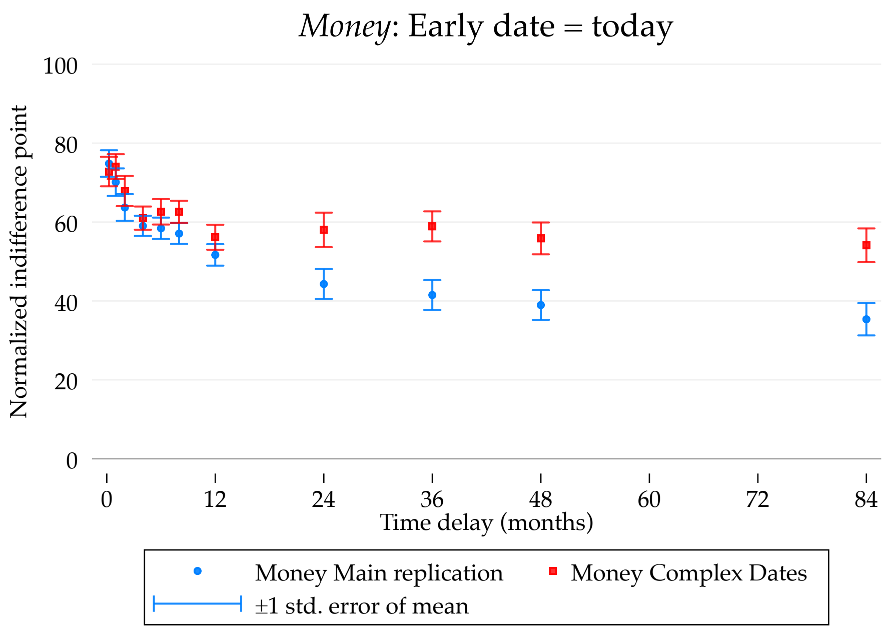

Figure 22: Observed discounting with $t_1 = 0$ in *Money Main replication* ($N = 161$ subjects) and *Money Complex Dates* ($N = 149$ subjects). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

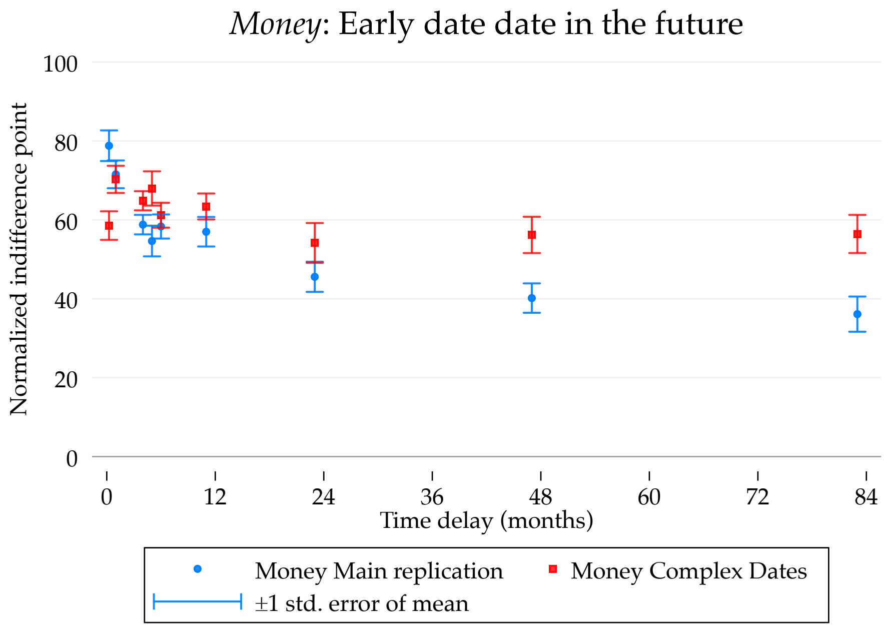

Figure 23: Observed discounting with $t_1 > 0$ in *Money Main replication* ($N = 161$ subjects) and *Money Complex Dates* ($N = 149$ subjects). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

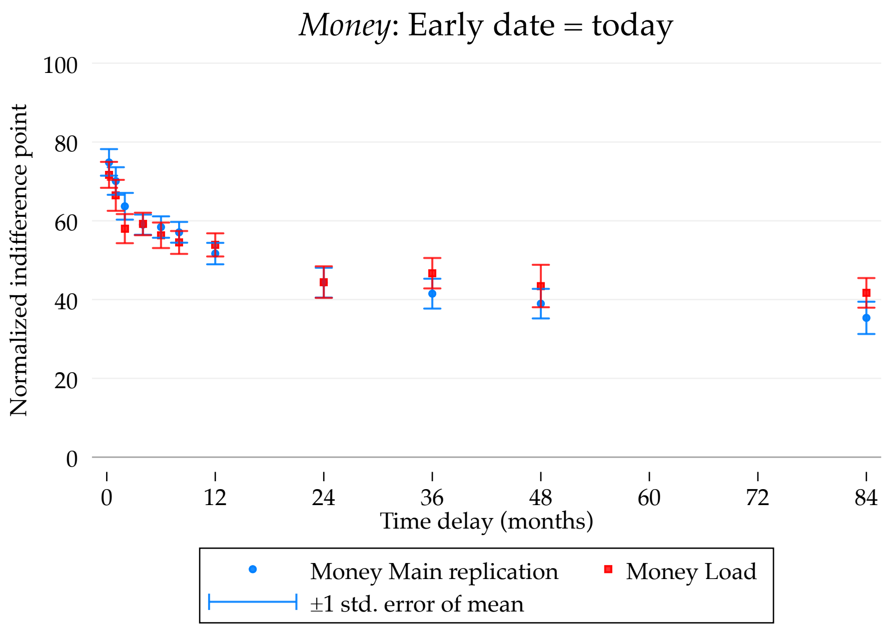

Figure 24: Observed discounting with $t_1 = 0$ in *Money Main replication* ($N = 161$ subjects) and *Money Complex Dates* ($N = 154$ subjects). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

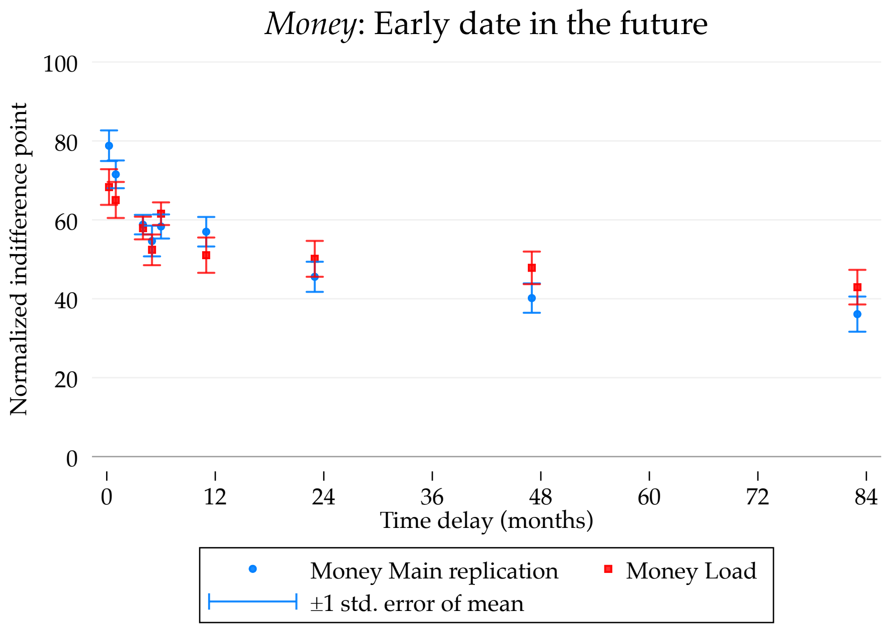

Figure 25: Observed discounting with $t_1 > 0$ in *Money Main replication* ($N = 161$ subjects) and *Money Complex Dates* ($N = 154$ subjects). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

F Results for Within-Subject Cognitive Load Experiment

In addition to the between-subject manipulation of cognitive load presented in Appendix E, we ran a within-subject version of this experiment. As in *Money Main*, subjects completed 12 choice lists involving hypothetical monetary payments. 6 of those rounds were NOLOAD rounds which were exactly identical to *Money Main*. In the other 6, randomly selected LOAD rounds, subjects faced a non-binding time limit of 15 seconds to complete the choice list and were instructed to sum up numbers that were flashed on the choice list screen in random intervals. There was no time limit and not number-counting task on the subsequent cognitive uncertainty elicitation screen. Following their cognitive uncertainty statement, subjects were prompted to enter the sum of numbers that we flashed on the choice list. Similar to treatment *Money Load*, our prediction was that in LOAD rounds subjects have relatively fewer cognitive resources available to fill out the choice list than in NOLOAD rounds. The manipulation should therefore lead to more pronounced insensitivity to time delays in LOAD rounds.

The predictions and sample size for this treatment were pre-registered at `https://aspredicted.org/av2y2.pdf`. We find that relative to NOLOAD rounds, the cognitive load manipulation increases average stated cognitive uncertainty by 9.2 percentage points (41%). Table 18 summarizes the effect of this treatment on intertemporal choices. They are in line with our predictions. In rounds with cognitive load, observed choices display substantially reduced sensitivity to the length of the time delay. These findings confirm our results from treatment *Money Load* at the subject level and suggest that the availability of cognitive resources is a source of cognitive imprecision that leads to more compressed intertemporal choice behavior.

Table 18: Cognitive load and insensitivity: Within-subject evidence

	Dependent variable: Normalized indifference point			
Sample:	$t1 = 0$		$t1 > 0$	
	(1)	(2)	(3)	(4)
Time delay (years)	-5.26***	-5.27***	-6.37***	-6.45***
	(0.43)	(0.43)	(0.50)	(0.50)
Time delay × 1 if cognitive load	2.33***	2.37***	3.81***	3.90***
	(0.63)	(0.63)	(0.70)	(0.71)
1 if cognitive load	-1.14	-1.15	-7.90***	-7.99***
	(1.03)	(1.03)	(1.33)	(1.33)
Payment amount FE	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes
Subject FE	Yes	Yes	Yes	Yes
Observations	2894	2894	1966	1966
R^2	0.70	0.70	0.68	0.68

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes decisions from the within-subject cognitive load experiment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

G Experimental Instructions

G.1 Money Main

Part 1 of this study: Instructions (1/3)

Please read these instructions carefully. There will be comprehension checks. If you fail these checks, you will immediately be excluded from the study and you will not receive the completion payment.

In this part of the study, you will **choose between various hypothetical payments, which pay different amounts of money at different points in time**. An example decision is between the following two hypothetical payments:

In 30 days: $ 40 OR **Today: $ 12**

For all hypothetical payments in this study, please treat them as if you knew that you would receive them with certainty, even if they are delayed. That is, please assume that there is no risk that you wouldn't actually get paid. Further assume that all payments were made by leaving a check in your mailbox.

Throughout the experiment, there are no right or wrong answers, because how much you like an option depends on your personal taste. There will be two types of decision screens.

Decision screen 1

On decision screen 1, you will be asked to choose which of two payment options you prefer. You will see choice lists such as the one below, where each row is a separate choice. In every list, the left-hand side option (Option A) is a delayed payment that is identical in all rows. The right-hand side option (Option B) is a payment *with an earlier payment date than Option A*. The earlier, right-hand side payment increases as you go down the list. **An effective way to complete these choice lists is to determine in which row you would like to switch from preferring Option A to preferring Option B.**

Based on where you switch from Option A to Option B in this list, we assess which amount at the early payment date (Option B) you value as much as the amount specified at the later payment date (Option A). For example, in the choice list below, you would value $40 in 30 days somewhere between $12 and $14 today, because this is where switching occurs.

Option A			Option B
In 30 days: $40	●	○	Today: $2
	●	○	Today: $4
	●	○	Today: $6
	●	○	Today: $8
	●	○	Today: $10
	●	○	Today: $12
	○	●	Today: $14
	○	●	Today: $16
	○	●	Today: $18
	○	●	Today: $20
	○	●	Today: $22
	○	●	Today: $24
	○	●	Today: $26
	○	●	Today: $28
	○	●	Today: $30
	○	●	Today: $32
	○	●	Today: $34
	○	●	Today: $36
	○	●	Today: $38
	○	●	Today: $40

On the next page, you will see an example choice list, and you can practice making your selections.

Click "Next" to proceed to the example page.

Part 1 of this study: Instructions (2/3)

Auto-completion: The table auto-completes your choices so you don't have to click through all of the rows. You do not have to start at the top of the table. If you select Option A in any one row, we assume that you will also prefer Option A in all rows *above* that row. If you select Option B in any one row, we assume that you will also prefer Option B in all rows *below* that row.

Option A			Option B
	○	○	**Today: $2**
	○	○	**Today: $4**
	○	○	**Today: $6**
	○	○	**Today: $8**
	○	○	**Today: $10**
	○	○	**Today: $12**
	○	○	**Today: $14**
	○	○	**Today: $16**
	○	○	**Today: $18**
In 30 days: $40	○	○	**Today: $20**
	○	○	**Today: $22**
	○	○	**Today: $24**
	○	○	**Today: $26**
	○	○	**Today: $28**
	○	○	**Today: $30**
	○	○	**Today: $32**
	○	○	**Today: $34**
	○	○	**Today: $36**
	○	○	**Today: $38**
	○	○	**Today: $40**

Part 1 of this study: Instructions (3/3)

Decision screen 2

When you fill out a choice list, you may feel **uncertain about whether you prefer the left or right payment option**. On decision screen 2, we will ask you to select a button to indicate **how certain** you are how much money the larger later payment is worth to you in terms of dollars at the earlier payment date.

In answering this question, we ask you to assume that you would receive both payment options with certainty. We are interested in **your uncertainty about your own preferences regarding these payments**, not in your potential uncertainty about whether you would actually receive the money.

Example

Suppose that on the first decision screen you indicated that you valued $40 in 30 days somewhere between $12 and $14 today. Your second decision screen would look like this.

How certain are you that you actually value $40 **in 30 days** somewhere between $12 and $14 **today**?

○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
0%	5%	10%	15%	20%	25%	30%	35%	40%	45%	50%	55%	60%	65%	70%	75%	80%	85%	90%	95%	100%

very uncertain **completely certain**

Comprehension questions

The questions below test your understanding of the instructions.

Important: If you fail to answer any one of these questions correctly, you will not be allowed to participate in the study.

1. Which of the following statements is true?

   ○ In making my decisions, I am asked to assume that I will actually receive all payments as indicated, regardless of whether they take place now or in the future.

   ○ In making my decisions, I am asked to assume that it is less likely that I will actually receive payments that are meant to take place in the future.

   ○ In making my decisions, I am asked to assume that it is less likely that I will actually receive payments that are meant to take place now.

2. Suppose you are 80% certain that your decisions actually correspond to how much the different choice options are worth to you. Which button should you click in this case?

○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
0%	5%	10%	15%	20%	25%	30%	35%	40%	45%	50%	55%	60%	65%	70%	75%	80%	85%	90%	95%	100%

very uncertain **completely certain**

3. When we ask you how certain you are about how much different payments are worth to you at different points in time, then which type of uncertainty are we interested in?

   ○ Uncertainty about whether I would actually receive the payments.

   ○ Uncertainty about how much I value the payments, assuming that I know I would receive them with certainty.

G.2 Voucher Main

Part 1 of this study: Instructions (1/4)

Please read these instructions carefully. There will be comprehension checks. If you fail these checks, we will have to exclude you from the study and you will not receive the completion payment.

In this part of the study, you will **choose between different UberEats food delivery vouchers. These vouchers will vary along two dimensions**:

- **The vouchers will have different values**
- **The vouchers will be valid at different points in time**

How do the vouchers work?

Each voucher is valid for food delivery during a period of only seven days. A voucher can be used starting **from the indicated date**, and **it remains valid for exactly 7 days after** that date. Specifically, the vouchers work as follows:

- If you win a voucher, you will be informed about the voucher amount and the validity period on the last page of this study. You will then be asked to provide an email address associated with an UberEats account. The voucher will directly be credited to the corresponding UberEats account within the next 10 hours.
- However, the voucher amount **can only be spent during the validity period** of the voucher.
- Vouchers can be used to order from the entire range of restaurants, cafes, and bars that partner with UberEats in your area.
- You do not need to worry about forgetting the validity period: **UberEats will automatically send reminders** about your voucher 24 hours before the validity period starts and 24 hours before it ends.

What decisions will you be asked to make?

An example decision is between the following two vouchers:

Valid in 30 days: $40 Voucher	OR	**Valid today: $20 Voucher**

The left-hand side voucher carries an amount of $40 and can be spent in the 7-day period starting in 30 days from now. The right-hand side voucher is for an amount of only $20, but can be spent in the 7-day period starting immediately.

Throughout the experiment, there are no right or wrong answers because how much you like a voucher depends on your personal taste.

Part 1 of this study: Instructions (2/4)

Decision screen 1

On decision screen 1, you will be asked to choose which of two vouchers you prefer. You will see choice lists such as the one below, where each row is a separate choice. In every list, the left-hand side option (Voucher A) is a voucher that is identical in all rows. The right-hand side option (Voucher B) is a voucher *with an earlier validity period than Voucher A*. The amount associated with the earlier, right-hand side voucher increases as you go down the list. **An effective way to complete these choice lists is to determine in which row you would like to switch from preferring Voucher A to preferring Voucher B.**

Based on where you switch from Voucher A to Voucher B in this list, we assess which voucher amount in the early validity period (Voucher B) you value as much as the voucher amount specified in the later validity period (Voucher A). For example, in the choice list below, you would value a $40 voucher that is valid in 30 days somewhere between a $12 and a $14 voucher that is valid today, because this is where switching occurs.

Voucher A			Voucher B
	●	○	Valid Today: $2 Voucher
	●	○	Valid Today: $4 Voucher
	●	○	Valid Today: $6 Voucher
	●	○	Valid Today: $8 Voucher
	●	○	Valid Today: $10 Voucher
	●	○	Valid Today: $12 Voucher
	○	●	Valid Today: $14 Voucher
	○	●	Valid Today: $16 Voucher
	○	●	Valid Today: $18 Voucher
Valid In 30 days: $40 Voucher	○	●	Valid Today: $20 Voucher
	○	●	Valid Today: $22 Voucher
	○	●	Valid Today: $24 Voucher
	○	●	Valid Today: $26 Voucher
	○	●	Valid Today: $28 Voucher
	○	●	Valid Today: $30 Voucher
	○	●	Valid Today: $32 Voucher
	○	●	Valid Today: $34 Voucher
	○	●	Valid Today: $36 Voucher
	○	●	Valid Today: $38 Voucher
	○	●	Valid Today: $40 Voucher

If you are selected to receive an additional reward from part 1 of the study, your reward will be determined as follows:
Your choice in a randomly selected row of a randomly selected choice list determines the amount of your personal voucher. Each choice list and each row are equally likely to get selected.

Important:

- Your choices may matter for real money! If you are selected to receive a bonus, one of your choices will actually be implemented, and your decision will determine which type of voucher you receive.
- Since only one of your decisions will be randomly selected to count, you should consider each choice list independently of the others. There is no point in strategizing across decisions.

On the next page, you will see an example choice list, and you can practice making your selections.

Click "Next" to proceed to the example page.

Part 1 of this study: Instructions (3/4)

Auto-completion: The table auto-completes your choices so you don't have to click through all of the rows. You do not have to start at the top of the table. If you select Voucher A in any one row, we assume that you will also prefer Voucher A in all *above* that row. If you select Voucher B in any one row, we assume that you will also prefer Voucher B in all rows *below* that row.

Reminder: both vouchers are valid for 7 days starting on the day indicated for each voucher.

Voucher A			Voucher B
Valid in 30 days: **$40 Voucher**	○	○	Valid today: **$2 Voucher**
	○	○	Valid today: **$4 Voucher**
	○	○	Valid today: **$6 Voucher**
	○	○	Valid today: **$8 Voucher**
	○	○	Valid today: **$10 Voucher**
	○	○	Valid today: **$12 Voucher**
	○	○	Valid today: **$14 Voucher**
	○	○	Valid today: **$16 Voucher**
	○	○	Valid today: **$18 Voucher**
	○	○	Valid today: **$20 Voucher**
	○	○	Valid today: **$22 Voucher**
	○	○	Valid today: **$24 Voucher**
	○	○	Valid today: **$26 Voucher**
	○	○	Valid today: **$28 Voucher**
	○	○	Valid today: **$30 Voucher**
	○	○	Valid today: **$32 Voucher**
	○	○	Valid today: **$34 Voucher**
	○	○	Valid today: **$36 Voucher**
	○	○	Valid today: **$38 Voucher**
	○	○	Valid today: **$40 Voucher**

Part 1 of this study: Instructions (4/4)

Decision screen 2

When you fill out a choice list, you may feel **uncertain about whether you prefer the left or right voucher**. On decision screen 2, we will ask you to select a button to indicate **how certain** you are about how much the larger voucher amount with the later validity period is worth to you in terms of voucher credit that can be spent in the earlier validity period.

Example

Suppose that on the first decision screen you indicated that you value a $40 voucher that is valid in 30 days somewhere between a $12 and a $14 voucher that is valid today. Your second decision screen would look like this.

How certain are you that you actually value a $40 voucher that is valid **in 30 days** somewhere between a $12 and a $14 voucher that is valid **today**?

○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
0%	5%	10%	15%	20%	25%	30%	35%	40%	45%	50%	55%	60%	65%	70%	75%	80%	85%	90%	95%	100%

very uncertain **completely certain**

Comprehension questions

The questions below test your understanding of the instructions.

Important: If you fail to answer any one of these questions correctly, you will not be allowed to participate in the study, and you will not receive the completion payment.

1. Which of the following statements about the voucher below is true?

   Valid in 1 month: $30 Voucher

   - ○ This voucher can be used to order food starting from today until no later than 1 month.
   - ○ This voucher can be used to order food any time after 1 month. The validity period has no end date.
   - ○ This voucher can be used to order food in the 7-day period starting in 1 month.

2. Suppose you are 80% certain that your decisions actually correspond to how much the different voucher options are worth to you.

   Which button should you click in this case?

○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
0%	5%	10%	15%	20%	25%	30%	35%	40%	45%	50%	55%	60%	65%	70%	75%	80%	85%	90%	95%	100%

   very uncertain **completely certain**

3. Which of the following statements is true?

   - ○ Even if the validity period starts in the future, my voucher will be credited to my UberEats account shortly after the experiment. I do not have to remember the validity period because UberEats will send me reminders.
   - ○ If the validity period of the voucher starts in the future, I should expect to get my voucher credited to my UberEats account only shortly before the validity period starts. I have to memorize the validity period, otherwise I may forget to use the voucher amount. There is also some risk that I will not actually receive the voucher.