Noisy Cognition and Intertemporal Choice*

Benjamin Enke   Thomas Graeber

August 12, 2022

Abstract

The most widely-studied and robust deviations of intertemporal choice behavior from the exponential discounted utility paradigm are typically conceptualized as resulting from non-standard preferences. We experimentally study the hypothesis that many of these patterns are instead largely driven by noisy cognitive processes, which lead people to implicitly treat different time delays alike to some degree. In our experiments, we measure cognitive uncertainty, which captures people's awareness of their noisiness. In the data, cognitive uncertainty strongly predicts various core empirical regularities, such as why people often appear very impatient over short horizons, why per-period impatience is smaller over long than over short horizons, why discounting is hyperbolic even when the present is not involved, and why choices violate transitivity. An account of noisy cognition also makes two new predictions, which we test and confirm: discounting is more hyperbolic when a decision is more complex, and cognitive uncertainty is strongly predictive of following expert advice.

Keywords: Intertemporal choice, bounded rationality, noisy cognition, complexity

*For helpful comments and discussions we are grateful to Michael Grubb (a discussant), Anujit Chakraborty, Cary Frydman, Alex Imas, Salvatore Nunnari, Severine Toussaert and Michael Woodford. We thank Tsahi Halyo and Zheng Li for outstanding research assistance. Enke acknowledges funding from the Foundations of Human Behavior Initiative and Mind, Brain and Behavior at Harvard. Enke: Harvard University, Department of Economics, and NBER, enke@fas.harvard.edu; Graeber: Harvard Business School, tgraeber@hbs.edu.

1 Introduction

Many important economic decisions such as those related to education, investment and saving have an intertemporal component, which motivated an extensive research program on intertemporal decision-making. A main result of this program is a set of empirical regularities that are inconsistent with the exponential discounted utility paradigm. As highlighted by the recent reviews of Ericson and Laibson (2019), Cohen et al. (2020) and Lipman and Pesendorfer (2013), a famous regularity is that people behave more impatiently in decisions that affect outcomes in the present than they do in decisions that only involve the future (sometimes paired with commitment demand). In response to this observation, the literature developed a broad class of models that Ericson and Laibson (2019) summarize as "present-focused preferences," which include models of present-biased preferences (Laibson, 1997) and temptation (Gul and Pesendorfer, 2001; Dekel et al., 2009). However, Cohen et al. (2020) and Ericson and Laibson (2019) emphasize that there are other widely-studied and economically no less important key empirical regularities that cannot be neatly explained with present-focused preferences alone. These regularities pertain to how people treat time delays of different lengths.[1]

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. As a result, people's implied per-period impatience strongly decreases in the length of the delay (Thaler, 1981; Loewenstein and Prelec, 1992), a behavior that *looks like* people treat different delays as more similar than they really are. This extreme flattening out of observed discounting for long horizons is not predicted by models of present bias (Laibson, 1997). Second, as shown in Panel B of Figure 1, an inelasticity of discounting with respect to the delay is also found when the earlier date is in the future, again at odds with a pure account of present-focused preferences (Kable and Glimcher, 2010).[2] Third, experimental studies robustly identify a transitivity violation called subadditivity, according to which people appear more patient in tradeoffs over one long interval than in choices where that same interval is partitioned into two sub-intervals (Read, 2001). Fourth, in addition to the regularities of hyperbolicity and subadditivity highlighted by Cohen et al. (2020), the literature has repeatedly documented that an experimentally-induced decrease in the availability of cognitive resources typically makes people less patient over short horizons but more patient over long ones (Ebert, 2001; Ebert and Prelec, 2007; Deck and Jahedi, 2015; Imas et al., 2021).

These four regularities appear to share the commonality that they reflect an insensi-

[1]As discussed in footnote 4, Cohen et al. (2020) also highlight a separate set of phenomena that are related to the structure of payouts rather than the structure of time delays. We focus on the latter.

[2]The first two patterns are also jointly referred to as the common difference effect or strongly diminishing impatience (Chakraborty et al., 2020).

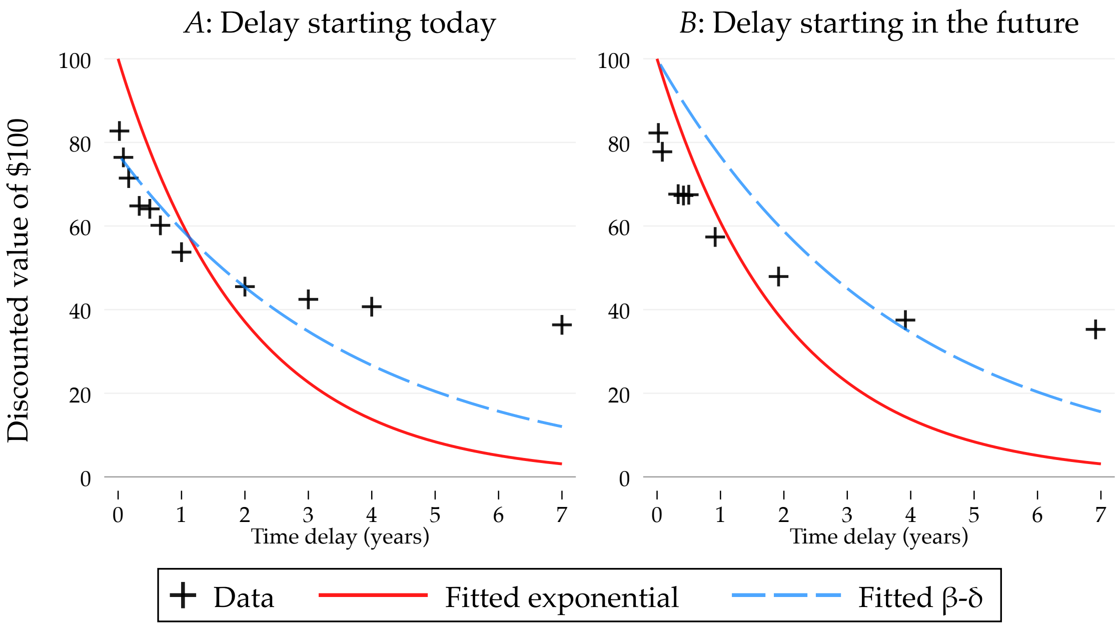

Figure 1: The figure shows the discounted value of $100 to be received with different time delays, partitioned by whether the early payment date is today (Panel A) or in the future (Panel B). The black markers indicate average behavior in our experiments described in Section 3.

tivity of decisions to the delay. A dominant approach in the behavioral economics literature has been to conceptualize this pattern by modifying people's discount functions. Either implicitly or explicitly, these approaches take the view that "anomalous" decisions are generated by "anomalous" preferences. However, as is well-known, the leading models that formalize non-standard preferences cannot account for all of the regularities summarized above. Moreover, the functional forms that generally fit experimental data best – the generalized hyperbola and its variants (e.g., Mazur, 1987; Loewenstein and Prelec, 1992; Kable and Glimcher, 2010) – were openly reverse-engineered to match patterns such as diminishing impatience, rather than developed because there was direct evidence for the existence of corresponding preferences. As a result, popular functional forms fit the data well in a reduced-form sense but do not provide micro-founded models of behavior, and do not explain why intertemporal decisions often vary as a function of cognitive states and environmental factors.

This paper experimentally studies the hypothesis that the empirical regularities, and intertemporal choice more generally, are to a large extent driven by noisy cognitive processes: people cognitively struggle with determining how much they value payments or consumption across different points in time. Our experimental analysis is organized around a discussion of theoretical models that explore how different versions of noise affect intertemporal decisions. For example, in Bayesian noisy cognition models (Woodford, 2020), the decision-maker may exhibit cognitive noise in determining how much a future payment is worth to him today. Such cognitive noise could have multiple origins, such as that the decision-maker does not know his true discount factor, that he has imperfect time perception, or that he struggles with integrating the time delays with his discount factor. In Bayesian noisy cognition models, the existence of cognitive noise induces the decision-maker to regress to (or anchor on) a cognitive default decision, such that the average observed decision is given by a convex combination of the true discounted-utility maximizing decision and the cognitive default. As a result, the decision-maker effectively treats different time delays as more similar to each other than they really are, which generates the aforementioned empirical regularities, including diminishing impatience and subadditivity. Similar predictions can emerge from a random response model in which the decision-maker probabilistically maximizes or chooses randomly. We discuss how random utility models and other cognitive noise models in the literature share some, but not all, of these implications.

Our main contribution is to experimentally test to what degree noisy cognition actually explains discounting behavior. To make the magnitude of cognitive noise visible, we leverage the insight that people often have some awareness of how noisy their decision process is. Specifically, our experiments measure people's *cognitive uncertainty*, which refers to a decision-maker's subjective uncertainty about their utility-maximizing decision (Enke and Graeber, 2022). Cognitive uncertainty captures a subject's composite awareness of cognitive noise in the determination of the decision and hence potentially includes a variety of aspects, such as uncertainty about one's discount factor or computational difficulties that arise in integrating the discount factor with the time delay.

We elicit intertemporal decisions in three different ways to show robustness. In a first paradigm, experimental participants make decisions in standard multiple price lists to trade off different UberEats vouchers that can be used for restaurant delivery and takeout. These vouchers are time-dated, such that actual consumption only occurs in a prespecified time period. In a second, complementary paradigm, we implement analogous decisions, except that these are defined over hypothetical monetary amounts. Third, we replicate our findings using a direct elicitation technique that does not rely on a visual price list representation of choice options. We discuss in detail how our study design relates to discussions about experimental intertemporal choice methodology, including reliability and fungibility of payments.

After each decision, we elicit cognitive uncertainty as a person's subjective probability that their stated valuation range of a larger-later payment (the switching interval in a choice list) actually contains their true valuation of the later payment. 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 validate our measure with incentivized choice data, by showing that that cognitive uncertainty is significantly correlated with across-trial variability in responses to repetitions of the same choice problem.

The main insight of our analysis, from which many of our results follow, is that cognitive uncertainty is strongly related to whether people's discounting looks like they treat different time delays alike to some degree. As a result of this compression effect, more cognitively uncertain decisions *look like* they reflect higher impatience over short horizons but lower impatience over very long ones (a "flipping" property).

Both the intensive and the extensive margin of cognitive uncertainty predict behavior. Indeed, the link between cognitive uncertainty and compression effects is strictly monotonic: people in the lowest cognitive uncertainty quartile respond more to time delays than people in the second quartile, who in turn respond more than those in the third quartile, and so on. This shows that the magnitude of cognitive uncertainty contains much information even away from the rational benchmark of zero, and that strictly positive cognitive uncertainty is not just driven by random measurement error.

Because cognitive uncertainty predicts to what extent people's decisions look like they treat different delays alike, it is strongly predictive of the empirical regularities mentioned in the motivating discussion: (i) short-run impatience and decreasing impatience (hyperbolicity) when the present is involved; (ii) short-run impatience and hyperbolic discounting when the present is not involved (such that notions of present bias do not apply); and (iii) the transitivity violation of subadditivity. All of these correlations are quantitatively large. For instance, the magnitude of decreasing impatience is five times larger with positive as opposed to zero cognitive uncertainty.

Our motivating hypothesis was that cognitive noisiness leads to an inelasticity of decisions with respect to the *length of the time delay*. Thus, as a placebo exercise, we pre-registered the prediction that cognitive uncertainty is *unrelated* to front-end delay effects, which refer to the pattern that people tend to be more impatient about a given delay that begins now rather than in the future. This prediction directly follows from our main hypothesis (and a simple model) because the length of a time delay is held constant in front-end delay experiments. In the data, we indeed see that cognitive uncertainty and front-end delay effects are uncorrelated. These results suggest that cognitive noise and present bias are complementary objects that explain different phenomena.

To complement our correlational analysis that rests on the measurement of cognitive uncertainty, we implement additional treatments in which we exogenously manipulate cognitive noise and trace effects on discounting behavior. To do so, we increase the complexity of the choice tasks by embedding a math problem into them. This complexity manipulation hones in on one specific source of cognitive noise among the different ones mentioned above. As predicted, we find that increased complexity (i) significantly increases cognitive uncertainty and (ii) leads to substantially more pronounced hyperbolic discounting. We interpret these patterns as saying that the hyperbolicity of observed discounting strongly depends on noisy cognition, which in turn partly depends on complexity.

To examine the quantitative importance of noisy cognition for predicting behavior, we estimate a cognitive noise model. In these estimations, allowing for cognitive uncertainty produces an increase in model fit that is quantitatively substantial. For example, relative to an exponential discounted utility model, the increase in fit is twice as large as the increase resulting from allowing for present bias.

All experiments summarized so far examine people's decision making when they are *forced* into making a decision. Yet, an account of cognitive noise in combination with awareness thereof (cognitive uncertainty) naturally also predicts that people might prefer not to make a decision themselves and instead follow expert advice. In contrast, in pure preferences-based accounts, people may behave in impatient ways, but they do not worry that the decision reflects a mistake. We find that cognitively uncertain participants are twice as likely to revise a previously-taken decision to follow the advice of professional economists. We interpret these patterns as suggesting that an account of cognitive noise and cognitive uncertainty not only helps in understanding which decisions people take but also whether they are likely to take a decision on their own in the first place.

In summary, our experiments document that the measurement of noisy cognition through cognitive uncertainty (i) sheds light on various core intertemporal choice regularities; (ii) facilitates the test of new predictions about how decision complexity affects discounting; (iii) has potential implications for choice architecture (advice seeking); and (iv) clarifies the welfare-relevant point that extreme short-run impatience and other non-standard behaviors partly reflect cognitive limitations rather than stable preferences.

Linking these insights to the literature, we propose that there are two different classes of anomalies that are related to variation in time delays. A first is the canonical evidence on front-end delay effects and dynamic inconsistency that are thought to reflect temptation or present bias (e.g., Laibson, 1997; Gul and Pesendorfer, 2001; Dekel et al., 2009; Toussaert, 2018; Chakraborty, 2021). A second, and larger, class of anomalies is captured by the stylized fact that intertemporal decisions are insufficiently sensitive to variation in the delay, which generates extreme impatience both when the present is involved and when it is not; hyperbolicity both when the present is involved and when it is not; subadditivity; and cognitive load effects that strongly depend on the length of the delay. A classic approach in the literature has been to develop models that capture both classes of phenomena. However, our evidence suggests that the two sets of regularities are driven by different principles: the second one is largely generated by noisy cognition, while the first one is not. Yet, noisy cognition 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.[3]

[3]These insights also bear an interesting relationship to Carrera et al. (2022) who identify noise as a

While ours is the first paper to empirically measure (awareness of) cognitive noise and show how it predicts discounting behavior, our approach builds on two different classes of models of cognitive or decision noise. A first class comprises the recent Bayesian cognitive noise literature (Woodford, 2020; Khaw et al., 2021; Gabaix, 2019; Frydman and Jin, 2021; Frydman and Nunnari, 2021). Gabaix and Laibson (2022), Gershman and Bhui (2019) and Vieider (2021) apply these models to intertemporal decisions. These papers are almost entirely theoretical, while we provide direct empirical evidence, both by measuring cognitive uncertainty and by experimentally manipulating cognitive noise. Moreover, in contrast to our focus on how people process time delays, Gabaix and Laibson (2022) model the noisy cognitive processing of utils. As a result, their model does not generate some of the key regularities that we are interested in (subadditivity) or does so only under fairly strong assumptions on prior beliefs (diminishing impatience).[4]

A second related class of models are random response and random preference models, which have received attention in decision theory and mathematical psychology (e.g., Lu and Saito, 2018; He et al., 2019). As we flesh out below, these models sometimes make predictions that are identical to Bayesian cognitive noise models. Relative to this literature, our contribution is to measure and exogenously manipulate cognitive noise, which allows us to provide much sharper and more direct tests than the model-fitting exercises that pervade the psychology literature on random choice (see Regenwetter et al., 2018, for a review).

The idea of empirically measuring cognitive noise and related concepts is increasingly gaining traction in the economics literature (e.g., Butler and Loomes, 2007; Agranov and Ortoleva, 2017; Khaw et al., 2021; Enke and Graeber, 2022).[5] Relative to Enke and Graeber (2022), where we document a link between cognitive uncertainty and probability weighting and belief formation, we here study an unrelated class of phenomena that (i) involve intertemporal decisions and (ii) differ from the S-shaped anomalies commonly identified in contexts involving probabilities. 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).

The paper proceeds as follows. Section 2 discusses theoretical background and de-

driver of commitment demand. Chakraborty et al. (2017) suggest a role for noise in driving present bias.

[4] The theoretical difference between time delays and utils also appears relevant because the empirical intertemporal choice regularities summarized in Cohen et al. (2020) can likewise be partitioned into those that concern time delays (such as subadditivity and hyperbolicity) and those that concern payouts or utils (gain-loss asymmetries and magnitude effects). While we focus on the former, the experiments in Gershman and Bhui (2019) suggest that magnitude effects may also be driven by cognitive noise.

[5] 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). Probably the biggest difference to our paper is that they do not focus on our objects of interest: explaining empirical regularities, studying the effects of complexity, and highlighting implications for advice seeking.

velops our predictions. Section 3 presents the experimental design and Sections 4–5 the results. Section 6 presents manipulations of complexity, Section 7 shows model estimations and Section 8 reports on our findings on advice following. Section 9 concludes.

2 Theoretical Considerations and Hypotheses

Consider a choice context in which a decision-maker (DM) is prompted to specify the units of consumption a in t_1 that make him indifferent to consuming $c_{t_2} = 1$ 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. A helpful theoretical benchmark is that of a rational DM's utility-maximizing decision, which equates the discounted utilities of both options.[6] 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 interest will be in how people's observed decision, denoted a^o, systematically deviates from a^* as a result of different types of cognitive noise. We only provide a brief discussion here because, as reviewed by Regenwetter et al. (2018), models that feature noise exhibit large diversity in precise modeling approaches and functional form assumptions. To begin, we will focus on two classes of models that – under certain assumptions – predict that the average observed decision can be represented as

$$E[a^o] = \lambda a^*(\delta, \Delta t) + (1-\lambda)d, \qquad \lambda \in [0,1], \qquad d \in [0,1], \tag{2}$$

where λ and d are explained below.

Bayesian cognitive noise models. We here outline a particular variant of Bayesian cognitive noise models that we use to derive our predictions. We then contrast it with other cognitive noise models in the literature and discuss to what extent they deliver similar predictions. Details are in Appendix A.

Many Bayesian cognitive noise (also called cognitive imprecision) models presume that the DM perceives some specific input parameter of the problem with noise (Woodford, 2020; Khaw et al., 2021; Frydman and Jin, 2021; Vieider, 2021). This is often interpreted as perceptual noise. As in Enke and Graeber (2022), we here take a broader perspective and formalize the idea that the DM exhibits cognitive noise in determining his utility-maximizing action, a^*. One interpretation of this is that people only have

[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.

access to a noisy cognitive representation of the relative magnitude of the effective discount factors of the two dates, $\delta^{\Delta t}$, which determines a^*. According to this broader perspective of cognitive noise, noisiness could result from any of the following: preference uncertainty about the discount factor; imperfect time perception (e.g., Zauberman et al., 2009; Brocas et al., 2018); and / or the cognitive difficulty of integrating the delay with the discount factor.

Formally, suppose the DM holds a Beta-distributed prior A over his discounted-utility maximizing action, where $A \sim Beta(n_1 d\ ,\ n_1(1-d))$.[7] The parameter n_1 reflects the DM's precision of his prior. We refer to the prior mean d as the cognitive default decision, which is the action that the DM would take before deliberating about the problem. The cognitive default reflects a *relative* valuation for a later option in terms of an earlier option (in percent). Crucially, the default is defined to be independent of the time delay in a specific choice problem. One way of thinking about this assumption is that the pre-deliberative default reflects the decision the DM would take before s/he has seen or parsed the delay in a specific problem.

Through deliberation, the DM generates a cognitive signal about what his discounted-utility maximizing decision a^* is. This signal S is (scaled) Binomially distributed, $S \sim \frac{1}{n_2} Bin(n_2, a^*)$, and is unbiased, $E[S] = a^*$. The parameter n_2 controls the precision of the cognitive signal. The subjective likelihood of the utility-maximizing action formed by a Bayesian DM based on a randomly drawn internal representation $\{S = s\}$ can then be represented by a binomial distribution:

$$\mathscr{L}(a^*|S=s) = P(S=s|a^*, n_2) = \binom{n_2}{sn_2}(a^*)^{sn_2}(1-a^*)^{(1-s)n_2}. \tag{3}$$

The posterior mean over the utility-maximizing action writes

$$a^o = \lambda(n_2)\cdot s + [1-\lambda(n_2)]\cdot d \ \Rightarrow\ E[a^o] = \lambda(n_2)\cdot a^* + [1-\lambda(n_2)]\cdot d. \tag{4}$$

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

In the formulation above, the signal precision is exogenous and constant. In our experimental data, we will see that cognitive noise is an increasing, strongly concave

[7] Our exposition is an adaptation of atemporal applications (Fennell and Baddeley, 2012; Heng et al., 2020; Enke and Graeber, 2022).

function of the length of the delay. Thus, in Appendix A we show that the predictions below also hold when λ exponentially decreases in Δt.

Pre-registered predictions. *Relative to a DM without cognitive noise ($\lambda = 1$), a DM with cognitive noise ($\lambda < 1$) 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.*

To see the logic behind these predictions, it is useful to imagine that the cognitive default action d in eq. (2) is "intermediate" in nature.[8] 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.

In our experiments, a central tendency for the mean prior belief is plausible because people have little experience with the context, akin to documentations of central tendency or compromise effects in psychology.[9] In our model estimations reported below, estimated defaults are indeed consistently intermediate.

The stylized model presented above differs in two respects from other prominent discounting applications of Bayesian noisy cognition in the literature, Gabaix and Laibson (2022) and Gershman and Bhui (2019).[10] First, both assume that all decision-relevant cognitive noise stems from the mental simulation of *future monetary amounts or utils*, while we posit noisy cognition at least partly with respect to effective discount factors. The main motivation is our objective to explain intertemporal choice patterns associated with variation in time delays (rather than, e.g., magnitude effects and gain-loss asymmetries). Second, our approach effectively assumes that DMs encode with noise the

[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]We experimentally distinguish this tendency towards an intermediate relative valuation from the potential heuristic of "clicking in the middle" of a multiple price list, see Section 5.4.

[10]In work that is contemporaneous to ours, Vieider (2021) models a DM who exhibits cognitive noise in the perception of time, combined with the assumption that time is encoded in log space. This model differs from ours in that it predicts that cognitive noise generates present bias, which we don't find empirically. Moreover, unlike us, Vieider neither directly measures nor experimentally manipulates cognitive noise or cognitive uncertainty. Rather, he fits his model to experimental data, as is typical in the psychology literature on random choice.

relative effective discount factors that are associated with the two delays, $\delta^{\Delta t}$. This is in contrast to Gabaix and Laibson, who assume that each option is represented independently before noisy representations are compared to one another. The assumption that discounting is relative (set-dependent) builds both on a theoretical literature (e.g., Ok and Masatlioglu, 2007; Read, 2001) and on studies that show that across different domains people make decisions by comparing dimension-by-dimension (Rubinstein, 2003; Arieli et al., 2011).

The distinctions between the framework above and Gabaix and Laibson (2022) matter. First, because Gabaix and Laibson (2022)'s model maintains transitivity, it does not predict subadditivity.[11] Second, their model only predicts diminishing impatience under fairly special assumptions on the signal structure and the location of the prior. Third, unlike ours, their model predicts that cognitive noise produces front-end delay effects and related preference reversals, see Section 2.7 of Gabaix and Laibson (2022).[12]

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

$$a^{tr} = \begin{cases} a^*(\delta, \Delta t) & \text{with prob. } \lambda \\ \epsilon & \text{otherwise} \end{cases} \qquad \Rightarrow \qquad E[a^{tr}] = \lambda a^*(\delta, \Delta t) + (1-\lambda)d. \tag{5}$$

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.

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 intertemporal decision is bounded by zero and one, random decision errors or measurement error (Gillen et al., 2019) may lead to boundary effects that push decisions to be more intermediate than they really are. 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 the boundaries (less than 5% of all decisions), and we have verified that identical results hold when these decisions

[11] Note that because our model predicts non-transitive behavior, the characterization of time consistency through stationarity and time invariance suggested by Halevy (2015) does not apply here.

[12] 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 $z > 0$.

are excluded from the analysis.

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

$$a^r = a^*(\tilde{\delta}, \Delta t). \tag{6}$$

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 (Weitzman, 2001; Jackson and Yariv, 2014; 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.

Summary and empirical implementation. The different classes of random choice models can make similar predictions. Moreover, the models afford varying degrees of flexibility (see Regenwetter et al., 2018). Hence, our objective is not to definitely tease 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 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 (2022), 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:

$$p_{CU} \equiv P(\left|A|\{S = s\} - a^o\right| > c), \tag{7}$$

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.

3 Experimental Design

3.1 Choice Tasks

Incentivized UberEats Voucher Experiments. In treatment *Voucher Main*, rewards are given by UberEats food delivery vouchers.[13] 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 7 shows a screenshot.

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.

[13]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. Thus, in these setups, we would be able to identify a potential link of cognitive uncertainty with short-run impatience but not with diminishing impatience.

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 feasible 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 5.4. 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 model estimations (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.

3.2 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 8 provides a screenshot. This cognitive uncertainty measurement follows the same protocol as proposed in a revised version of Enke and Graeber (2022) for choice under risk, here adapted to an intertemporal choice context. In line with the discussion in Section 2, 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, in particular their discount factor. 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., Yeung and Summerfield, 2012; 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). 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.

Incentives. We deliberately do not incentivize the CU elicitation to maintain the simple – and for subjects intuitive – nature of the protocol. Recent research highlights that adding complex and potentially confusing scoring rules to elicitation tasks may actually decrease the quality of responses (Danz et al., 2022). To the degree that a lack of incentives may induce noisier CU data, our results on the link between CU and intertemporal decisions will be biased downward.

Link with choice data. 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 our unincentivized question, in the spirit of recent experimental validation studies in the literature (e.g. Falk et al., 2015; Enke et al., forthcoming). In both of our datasets (money and vouchers), across-trial variability and cognitive uncertainty exhibit a correlation of $r \approx 0.17$, $p < 0.01$ (Appendix Figure 10). This is in line with the correlations between cognitive uncertainty and choice variability in lottery choice and belief updating documented in Enke and Graeber (2022). Arts et al. (2020) also show correlations between randomization and an unincentivized confidence question. In combination, we interpret these results as strongly suggesting that incentivized choice data are consistent with the idea that our unincentivized cognitive uncertainty measurement captures the noisiness of people's cognitive processes.

3.3 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; Chakraborty et al., 2020). 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.

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.[14]

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 their own valuation, given that they know they will receive the money

[14]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.

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. 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 by utility curvature. 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.

3.4 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.[15] 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.[16] As we pre-registered, $N = 500$ workers participated in the UberEats study.

Participants in both studies completed a comprehension check quiz 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 E contains screenshots of all experimental instructions and comprehension checks.

3.5 Pre-Registration

Appendix Table 6 provides an overview of all treatments conducted for this paper, including pre-registration details. Our pre-registration includes (i) predictions 1–4 in Section 2, (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

[15] 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.

[16] 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.

discussed in Section 4.

4 Descriptives

Appendix Figure 9 shows histograms of task-level CU in the MPL decisions in treatments *Money Main* (left panel) and *Voucher Main* (right panel), such that each participant contributes twelve observations. 75% of all decisions in *Money Main* and 81% of decisions in *Voucher Main* are associated with strictly positive CU. This heterogeneity reflects both across-participant heterogeneity and systematic variation across choice problems. Figure 2 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 increases in the length of the absolute time delay up to a delay of about one year ($r = 0.07, p < 0.01$, for delays smaller than 24 months). This suggest that payouts or consumption in two temporally distant periods are generally more difficult to evaluate against each other. Strikingly, the observed pattern is strongly concave, and the relationship is essentially flat for delays longer than a year ($r = -0.05, p = 0.35$, for delays of 24 months or longer). In Appendix A, we show that our theoretical predictions apply when the magnitude of cognitive noise is a concave increasing function of the delay rather than a constant.

A relevant question is how consistently people exhibit 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 2 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*.[17]

5 Cognitive Uncertainty and Intertemporal Choice

5.1 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 indifference point*, which is given by the midpoint of the switching interval, divided by the

[17]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.

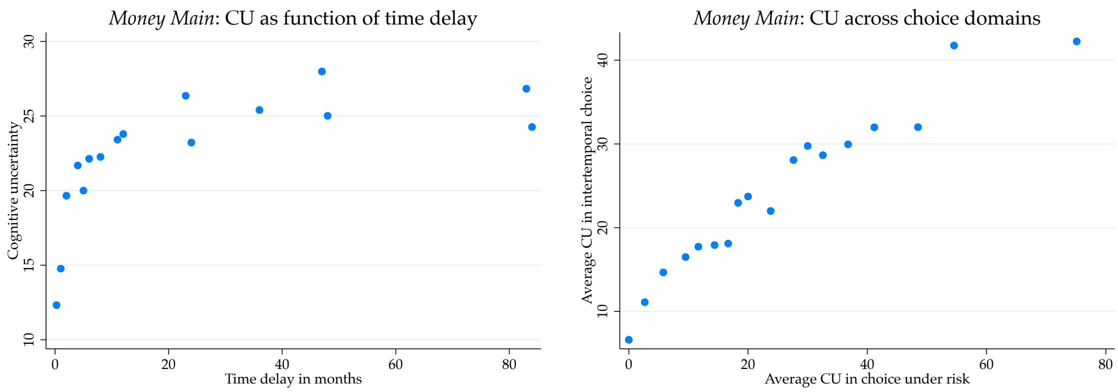

Figure 2: Binscatter plots. The left panel shows the relationship between task-level CU and the log time delay in a decision problem (N=7,740 decisions). The right panel shows the correlation between participant-level average CU in intertemporal choice and average CU in choice under risk (N=645 participants).

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 3 illustrates the relationship between normalized indifference points (in percent), cognitive uncertainty and time delays. The left-hand panels show results for treatment *Money Main* and the right-hand panels those for *Voucher Main*. In the top panels, we plot normalized indifferent points separately for participants with CU of zero and strictly positive CU. To make the results comparable between the voucher and money experiments, the x-axes are kept identical even though the maximal time delay in the vouchers study is only twelve months. 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 11 in Appendix B). The main takeaway of the top panels 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 cognitive noise in intertemporal choice is indeed insensitivity to time delays, rather than universally higher impatience. A second takeaway is that behavior is very similar in *Money Main* and *Voucher Main*, including in its link to CU. Finally, notice that observed normalized switching points tend to values below 50% for long delays. This is in line with our modeling framework sketched in Section 2 even under a default decision of 50%, see Section 7 for a discussion of this issue.

The bottom panels of Figure 3 provide a more complete picture of the relationship between cognitive uncertainty and sensitivity to time delays that does not rely on the arbitrary sample split into $CU = 0$ and $CU > 0$. We now split the sample into cognitive

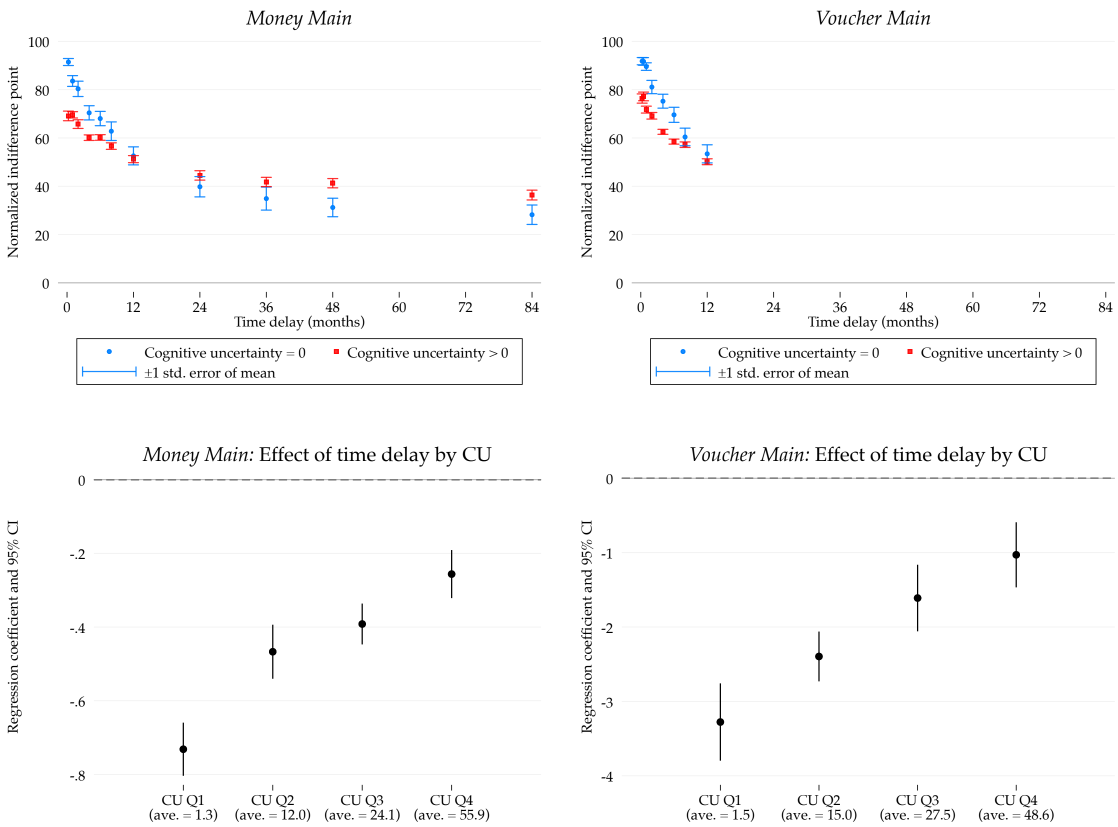

Figure 3: Top panels: Normalized switch points as a function of time delay with $t_1 = 0$ in *Money Main* (left, $N = 4,948$) and *Voucher Main* (right, $N = 3,846$). Normalized indifference points are given by the midpoint of the switching interval in a choice list, divided by the larger-later payout amount (in %). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level. Bottom panels: Coefficients from regressions of normalized indifference points on time delay, split by CU quartiles (for comparability again $t_1 = 0$ only; left: *Money Main*; right: *Voucher Main*).

uncertainty quartiles. Because in our experiments between 20% and 25% of all CU statements are equal to zero, the first quartile almost corresponds to $CU = 0$, while the other quartiles leverage variation in the intensive margin of CU. For each of the four CU buckets, we regress observed indifference points on the time delay and report the coefficient. If discounting did not depend on cognitive noise, the four regression coefficients would be equally large. Instead, we see that the effect of the time delay on behavior continuously decreases (in absolute terms) as CU increases. This shows that the results are not just driven by the extensive margin of CU, but that higher CU is strongly associated with more compression also within the sample of strictly positive CU.

Table 1 presents further supporting OLS regression estimates. Here, we relate participant'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 3. First, CU is associated

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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

with a lower sensitivity of indifference values with respect to time delays, as can be inferred 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 $t_1 > 0$. To illustrate 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]

5.2 Linking Cognitive Uncertainty to Empirical Regularities

5.2.1 Short-Run Impatience

Figure 3 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 correlations 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 regressions.

[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.

5.2.2 Decreasing Impatience

To study decreasing impatience,[19] 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 works with a per-period measure of patience as $RRR/\Delta t$. A transformation of this measure that captures per-period patience in an intuitive way is

$$\delta_H(a^o) \equiv e^{-RRR/\Delta t} = (a^o)^{1/\Delta t}. \tag{8}$$

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 4 shows the link between CU and decreasing impatience in four different panels: treatments *Money Main* and *Voucher Main*, separately for $t_1 = 0$ and $t_1 > 0$. For each sample, we compute the average implied $\delta_H(a^o)$ across subjects for a given delay.[20]

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. Furthermore, we have again verified that the results are very similar if we exclude decisions that are associated with $CU = 0$.

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 $t_1 > 0$, 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.

[19]Decreasing impatience is by far the dominant finding in the literature. However, it is not universal, neither when the early date is today nor when it is in the future (see, e.g., Harrison et al., 2005).

[20]This figure 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 4 should be constant in the delay.

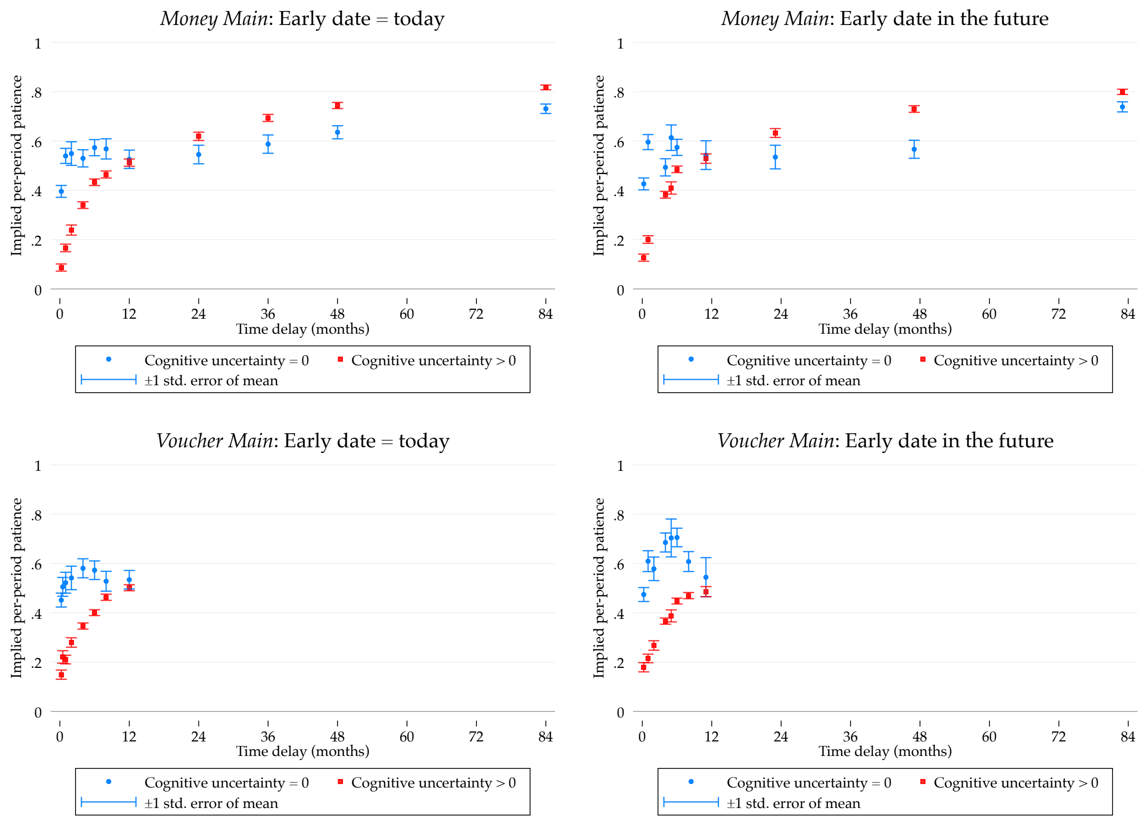

Figure 4: Implied per-period patience in *Money Main* (top panels) and *Voucher Main* (bottom panels), partitioned by whether the early payment date is today or in the future. Per-period patience is computed as $\delta_H(a^o) \equiv e^{-RRR/\Delta t} = (a^o)^{1/\Delta t}$, where a^o is the observed normalized indifference point. The figure shows average δ_H across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

5.2.3 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}. Following standard procedures in the literature, we compare the normalized indifference point obtained from the problem involving 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 (Read, 2001; Dohmen et al., 2017). Table 2 summarizes the results for both *Money Main* 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).

Table 2: Cognitive uncertainty and subadditivity

	Dependent variable: Normalized indifference point over long interval					
Treatment:	*Money Main*			*Voucher Main*		
	(1)	(2)	(3)	(4)	(5)	(6)
1 if 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}. 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. Demographic controls include age, gender and income. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2.4 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}. Because our main hypothesis is that people mentally distort the relative discount factors associated with two delays (or even the length of the delay itself), we predicted and pre-registered that cognitive uncertainty is uncorrelated with front-end delay effects as these experimental effects hold the length of the delay constant. 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) of Table 3 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 correlation between front-end delay effects and cognitive uncertainty is either small

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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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*.

5.2.5 Cognitive Load Effects

Using manipulations of time pressure or enforced waiting periods, the existing literature has documented that the availability of cognitive resources affects intertemporal choices (Deck and Jahedi, 2015; Imas et al., 2021; Ebert, 2001). While these studies typically implemented relatively short time delays and found higher impatience when the availability of cognitive resources is decreased, Ebert (2001) and Ebert and Prelec (2007) actually find that time pressure strongly affects the *hyperbolicity* of the discount function, such that for sufficiently long time horizons, time pressure actually leads to *higher patience*. Our framework predicts these patterns if and only if a decreased availability of cognitive resources increases cognitive noisiness.

To test the hypothesis that cognitive load increases both cognitive noise and the hyperbolicity of discounting, we designed the 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 ran-

Table 4: Effect of cognitive load and complexity manipulations on discounting

	Dependent variable: Normalized indifference point							
	Money Main Repl. vs. *Money Load*				*Money Main Repl.* vs. *Money Complex Dates*			
Sample:	$t1 = 0$		$t1 > 0$		$t1 = 0$		$t1 > 0$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Time delay (years)	-5.01***	-4.89***	-4.80***	-4.85***	-5.01***	-4.97***	-4.80***	-4.84***
	(0.55)	(0.55)	(0.63)	(0.62)	(0.55)	(0.55)	(0.63)	(0.62)
Time delay × 1 if *Load*	1.58**	1.64**	1.96**	2.00**				
	(0.78)	(0.78)	(0.82)	(0.82)				
1 if *Load*	-2.59	-2.34	-2.84	-2.68				
	(3.07)	(3.04)	(3.05)	(3.04)				
Time delay × 1 if *Complex Dates*					2.96***	2.97***	3.33***	3.36***
					(0.79)	(0.79)	(0.88)	(0.88)
1 if *Complex Dates*					2.93	3.17	1.06	1.38
					(2.99)	(3.00)	(2.94)	(2.94)
Payment amount FE	No	Yes	No	Yes	No	Yes	No	Yes
Demographic controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	2428	2428	1352	1352	2381	2381	1339	1339
R^2	0.07	0.10	0.06	0.07	0.07	0.08	0.06	0.07

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. Demographic controls include age, gender and income bucket. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

dom intervals next to the choice list.[21] An obvious issue with this load manipulation is that cognitive effort and resulting response times are endogenous: in principle, it is conceivable that subjects in the load condition 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 *Money Main* that we administered in the same experimental sessions. We conducted these experiments with a separate sample of 617 participants, in which each participant was randomly assigned to one of four treatments: *Money Load*, *Money Main Replication* and two further treatments discussed in Section 6.

We find, first, that our cognitive load manipulation increases stated CU relative to the replication of our main treatment by 5 percentage points (19%), $p < 0.01$. Second, we observe that cognitive load indeed leads to substantially more hyperbolic discounting, see columns (1)–(2) and (5)–(6) of Table 4. We interpret these patterns as suggesting that cognitive load increases the magnitude of cognitive noise, which in turn increases the hyperbolicity of discounting.

[21] 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 and reported in an earlier version of this paper (Enke and Graeber, 2021).

5.3 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 2 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 2 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 (2022), 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 (2022) and Gershman and Bhui (2019) do not explain all aspects of the evidence.

5.4 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. As we document in Appendix Table 9, the results remain statistically significant conditional on these subject fixed effects.

Measurement error. Our paper contributes to a literature on noise and measurement error in experiments (e.g., Gillen et al., 2019; Andersson et al., 2020). While we now discuss why measurement error in the sense of random response noise alone cannot explain our results, the measurement of cognitive uncertainty provides a useful complement to this literature because it allows researchers to predict *systematic* (rather than mean-zero) distortions of intertemporal decisions and preference estimates.

In principle, CU and intertemporal choices could be subject to a form of correlated experimental measurement error (Gillen et al., 2019) that would potentially create a mechanical relationship between the occurrence of strictly positive CU and the sensitivity of intertemporal decisions to delays. To illustrate, suppose that all subjects actually exhibit zero cognitive noise. Further suppose that (i) more inattentive subjects are more likely to exhibit random measurement error in the CU elicitation that leads them to state strictly positive CU, and (ii) that this same inattention will also lead subjects to make intertemporal decisions that are insensitive to time delays. Under this logic, CU and intertemporal decisions would be mechanically correlated. If this were the case, however, we would expect that CU has no predictive power for intertemporal decisions within the sample of strictly positive CU. As Figure 3 showed, this is counterfactual as the sensitivity of decisions to delays strongly decreases in CU, also conditional on CU > 0.

A second approach to investigate the importance of measurement error in the CU elicitation for our results is to look at subjects' across-task variability in CU, as presumably those subjects that are more inattentive and exhibit greater measurement error should make CU statements that are more variable. To study this, we compute the subject-level standard deviation of CU across experimental decisions, and study whether it is correlated with the sensitivity of decisions with respect to the delay, just like we did for the level of CU in Table 1. We find that the subject-level SD and sensitivity to delays are essentially uncorrelated ($p = 0.87$). This again speaks against the relevance of measurement error in CU for generating our main findings.

Direct elicitation experiments without price list format. Up to this point, all results were derived from experiments in which intertemporal choice behavior was elicited using choice lists. While this elicitation procedure is standard, it raises the potential concern that price lists have their own effects on behavior, in particular that they may induce subjects to switch around the middle of the list (e.g., Beauchamp et al., 2019). To document that the logic of CU and inelasticity extends to another elicitation technique, treatment *Money Main* also included a direct elicitation component that has no visual price list grid, see Section 3. 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 above: (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.

6 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. The effects of cognitive load (Section 5.2.5) already hint at the critical role of cognitive resources. We here conjecture that the magnitude of cognitive noise will also be affected of the exogenously determined complexity of the decision problem.

We manipulate the complexity of processing the time delay. 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."[22] As described in Section 5.2.5, the complexity experiments were conducted in joint sessions with the load manipulation and a replication of the baseline condition ($N = 617$ participants, random assignment to treatments). Here, too, we implemented a time limit of 25 seconds per choice. Our hypothesis is that the treatment increases cognitive noise in the determination of the utility-maximizing decision. While the treatment likely does not change uncertainty over one's true discount factor, it plausibly changes the difficulty of combining the discount factor with the time delay (recall that we embrace different potential sources of cognitive noise).

We find that the complexity variation substantially increases stated CU relative to the replication of our main treatment, by 12 percentage points (50%). As columns (5)–(8) of Table 4 show, the complexity manipulation also makes discounting decisions substantially less sensitive to variation in the time delay, which produces more pronounced hyperbolic discounting. These results are at odds with all models that formalize hyperbolicity as reflecting stable preferences. Instead, we interpret them as suggesting that higher complexity causes more pronounced cognitive noise, which leads to more pronounced hyperbolicity.

7 Model Estimations

We proceed by estimating eq. (2) from Section 2 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

[22]In another complexity manipulation, treatment *Money Complex Amounts*, we instead represented the monetary amount for choice option A as a math problem, such as "\$(4*8/2)+(8*9/2)-12". The results of this treatment are very similar to those for *Money Complex Dates*, see Appendix Table 10. For example, average CU increases by 10 percentage points (42%) in this treatment.

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

$$\begin{aligned} \mathbb{E}[a^o] &= \lambda(p_{CU}) \cdot \mathbb{E}[s] \cdot x + (1 - \lambda(p_{CU})) \cdot d \cdot x \\ &= (1 - \alpha \cdot p_{CU}) \cdot (\delta^{\Delta t})^{1/\gamma} \cdot x + (\alpha \cdot p_{CU}) \cdot d \cdot x \end{aligned} \tag{9}$$

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.[23] To assess and compare model fit, we estimate four model variants. First, a baseline exponential discounting model that ignores cognitive noise (i.e., we set $\alpha = 0$). Second, also for benchmarking purposes, a $\beta - \delta$ model, which also precludes a role for cognitive noise. Finally, we estimate both of these variants including CU.[24]

Our objective here is to pit CU-amended models against canonical benchmark models that have a direct psychological motivation, rather than to argue that there are no other functional forms that fit the data better. As discussed in the Introduction, functional forms such as the generalized hyperbola were in fact designed to fit observed data well, but are arguably not directly psychologically microfounded in the way that accounts of present-focused preferences or cognitive noise are.

Notice that, following the discussion in Section 2, 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 5 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,

[23]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.

[24]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 5: 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 (9). 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.

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.

Figure 5 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. The figures are constructed by generating predicted values, based on the parameter estimates in Table 5. We again see that exponential discounting fits the data poorly. Likewise, almost by construction, the canonical beta-delta model fits poorly when the early payment date is in the future.[25] 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 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. Note that even with an estimated cognitive default of around 50%, predicted indifference points in the delta-CU model do not converge to 50% for very long time horizons, but a significantly lower value, mirroring the convergence pattern in the actual data. The reason is that even though cognitive noise increases in the time horizon as suggested by Figure 2, this relationship is strongly concave and essentially flat for delays longer than 24 months. Thus, λ does not converge to one. As a consequence, decisions in the delta-CU model attribute strictly positive weight to

[25]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 $t_1 = 0$.

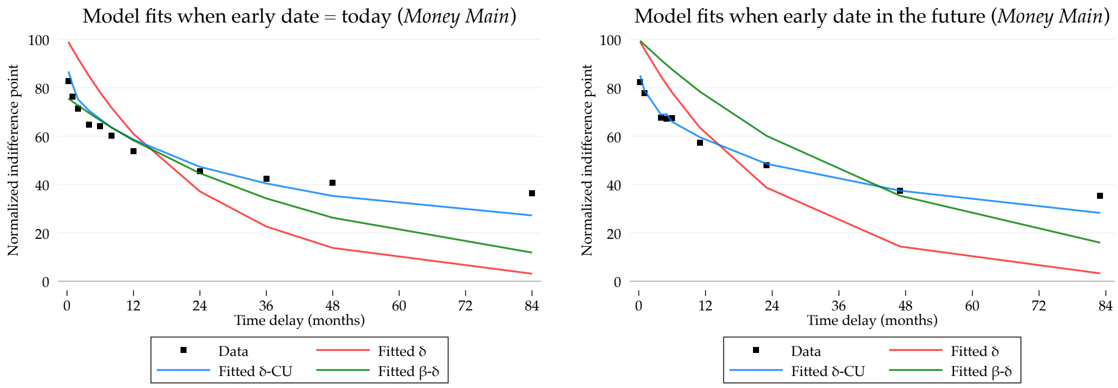

Figure 5: Model fit vs. data in *Money Main*. The model predictions are computed as fitted values of the parameter estimates in Table 5.

the discounted-utility maximizing action a^* even for very long horizons, which pulls observed indifference points below 50%.

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 model at the level of individual subjects.[26] We report the results in Appendix Table 11. To summarize, there is substantial individual-level variation in estimated parameters. For most parameters, the center of the estimated coefficient distributions is line with the parameters in our representative-agent estimation.[27]

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. 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. This effect is generally referred to as "central tendency effect" in psychology (Hollingworth, 1910), and has also been to cognitive noise (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).

[26]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.

[27]An exception is 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 quasi-hyperbolic discounting can partly result from the aggregation of individuals with heterogeneous discount factors.

8 Advice Following

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 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`.

Experiments 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.

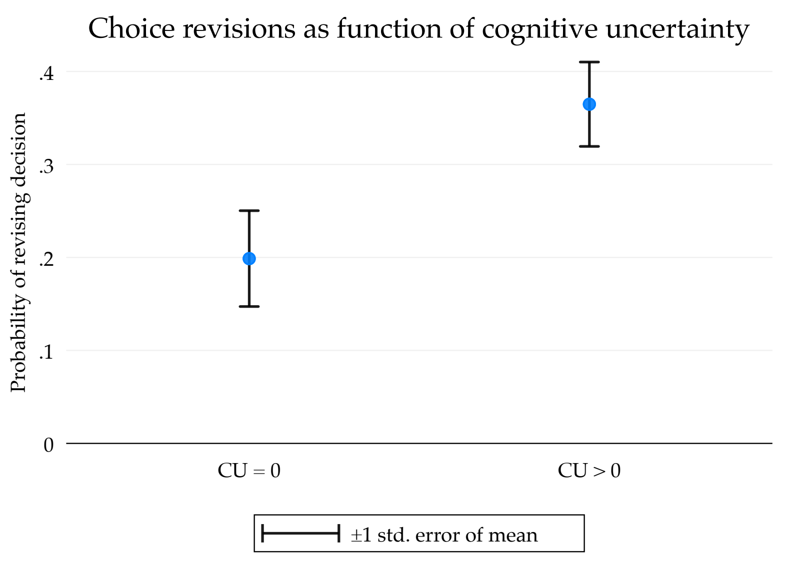

Figure 6: Probability of revising decision towards higher patience, as a function of cognitive uncertainty ($N = 153$). The figure is constructed controlling for the normalized indifference point before seeing advice. In other words, the y-axis shows the residual probability of revising the decision 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 6 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 and relevance. Much of behavioral economics views intertemporal choice and its 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. We document that a large share of short-run impatience, hyperbolic discounting and subadditivity effects are driven by bounded rationality and cognitive noise, rather than impatient

[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 6 from the initial normalized indifference point through a linear regression (the results are even stronger without this adjustment).

preferences.

Distinguishing whether intertemporal choice is only driven by non-standard preferences or also by cognitive limitations is important for at least three reasons. First, trivially, the welfare implications differ. Second, as we have shown, mistakes (and awareness thereof) generate systematic demand for expert advice or policy that is absent in preferences-based models. Third, unlike preferences-based models, accounts of cognitive noise predict that decisions should be strongly influenced by the complexity of the decision environment, which is practically relevant.

Limitations. As noted in the Introduction, our paper does not purport to explain all intertemporal choice anomalies. First, we focused on the most well-known and most robust regularities that relate to variation in the *time delay*, but ignored those regularities that pertain to *payout effects* (gain-loss asymmetries and magnitude effects). Second, our study does not address 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.

Links between literatures on risky and intertemporal choice. In classical models of decision-making, preferences over risk and intertemporal tradeoffs are captured by separate parameters. In practice, however, it is well-known that intertemporal decisions are a function of risk, for example that the presence of uncertainty can confound the experimental measurement of time preferences (e.g., Halevy, 2008). In conjunction with the results reported in Enke and Graeber (2022), the findings in this paper suggest that there may be a second deep link between intertemporal and risky decisions, which is that they are linked through the presence of cognitive noise. For example, if cognitive noise leads to both more extreme probability weighting and more pronounced diminishing impatience, then conventional experimental techniques to measure preference parameters will identify links between the two phenomena (as in, e.g., Epper et al., 2011). Another example concerns the well-known correlations between cognitive ability, risk taking and patience (e.g., Dohmen et al., 2010; Falk et al., 2018). If cognitive ability (as measured in experiments or surveys) partly picks up cognitive noise, then this may explain why empirical measures of risk and time preferences are correlated both with each other and with cognitive ability.

References

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

Andersson, Ola, Håkan J Holm, Jean-Robert Tyran, and Erik Wengström, "Robust inference in risk elicitation tasks," *Journal of Risk and Uncertainty*, 2020, *61* (3), 195–209.

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.

Arieli, Amos, Yaniv Ben-Ami, and Ariel Rubinstein, "Tracking decision makers under uncertainty," *American Economic Journal: Microeconomics*, 2011, *3* (4), 68–76.

Arts, Sara, Qiyan Ong, and Jianying Qiu, "Measuring subjective decision confidence," *Working Paper*, 2020.

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.

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," *The Review of Economic Studies*, 2022, *89* (3), 1205–1244.

Chakraborty, Anujit, "Present bias," *Econometrica*, 2021, *89* (4), 1921–1961.

— , **Evan M Calford, Guidon Fenig, and Yoram Halevy**, “External and internal consistency of choices made in convex time budgets,” *Experimental economics*, 2017, *20* (3), 687–706.

— , **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.

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

Danz, David, Lise Vesterlund, and Alistair J Wilson, “Belief elicitation and behavioral incentive compatibility,” *American Economic Review*, 2022.

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.

Dekel, Eddie, Barton L Lipman, and Aldo Rustichini, “Temptation-driven preferences,” *The Review of Economic Studies*, 2009, *76* (3), 937–971.

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.

— , — , — , **and** — , “The robustness and pervasiveness of sub-additivity in intertemporal choice,” in “Working Paper” 2017.

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 in intertemporal choice,” Technical Report, National Bureau of Economic Research 2021.

— **and** — , “Cognitive uncertainty,” Technical Report, National Bureau of Economic Research 2022.

— , **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.

__ , **Helga Fehr-Duda, and Adrian Bruhin**, "Viewing the future through a warped lens: Why uncertainty generates hyperbolic discounting," *Journal of risk and uncertainty*, 2011, *43* (3), 169–203.

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

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global evidence on economic preferences," *The Quarterly Journal of Economics*, 2018, *133* (4), 1645–1692.

__ , __ , __ , **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.

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 2022.

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

Gillen, Ben, Erik Snowberg, and Leeat Yariv, "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study," *Journal of Political Economy*, 2019.

Gul, Faruk and Wolfgang Pesendorfer, "Temptation and self-control," *Econometrica*, 2001, *69* (6), 1403–1435.

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.

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

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.

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

Lipman, Barton L and Wolfgang Pesendorfer, "Temptation," in "Advances in economics and econometrics: Tenth World Congress," Vol. 1 Citeseer 2013, pp. 243–288.

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.

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.

Ok, Efe A and Yusufcan Masatlioglu, "A theory of (relative) discounting," *Journal of Economic Theory*, 2007, *137* (1), 214–245.

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.

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.

Toussaert, Séverine, "Eliciting Temptation and Self-Control Through Menu Choices: A Lab Experiment," *Econometrica*, 2018, *86* (3), 859–889.

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

Vieider, Ferdinand M, "Noisy coding of time and reward discounting," Technical Report, Ghent University, Faculty of Economics and Business Administration 2021.

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.

Yeung, Nick and Christopher Summerfield, "Metacognition in human decision-making: confidence and error monitoring," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 2012, *367* (1594), 1310–1321.

Zauberman, Gal, B Kyu Kim, Selin A Malkoc, and James R Bettman, "Discounting time and time discounting: Subjective time perception and intertemporal preferences," *Journal of Marketing Research*, 2009, *46* (4), 543–556.

ONLINE APPENDIX

A Derivations for Bayesian Cognitive Noise Model

A.1 Model Setup

Below we derive the main behavioral predictions of the Bayesian cognitive noise model introduced in Section 2. Recall that the DM's observed action given a cognitive signal is assumed to be the posterior mean:[30]

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

A more precise mental simulation (higher n_2) has a 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. The characterization of cognitive uncertainty (and its relationship to cognitive noise) is identical to the one provided in the Appendix of Enke and Graeber (2022).

A central intuition about the determinants of cognitive noise in intertemporal decisions is that longer delays are harder to simulate. In Section 4, we confirmed that stated cognitive uncertainty is increasing and strongly concave in the length of the delay under consideration. To accommodate this, we derive our predictions for the general case where λ is either constant or a decreasing concave function of Δt. Specifically, we model an exponentially decaying relationship between cognitive noise and time delays:

$$\lambda := \lambda(t) = c + be^{-k\Delta t}. \tag{11}$$

with positive k and $c, b \in [0, 1]$. For $b = 0$, λ is constant in the time delay. In the proofs we will distinguish between different cases wherever necessary.

A.2 Derivations for Behavioral Predictions

All theorems and derivations in this subsection will be provided for a given subject's mean observed action, i.e., their average response aggregating across many unbiased

[30] We focus on the posterior mean mostly for tractability. The subjectively optimal action depends on the assumption about the implicit loss function. For example, quadratic loss would imply the mean, absolute loss would imply the median as the optimal choice. In principle, people's subjectively expected reward from a given response might also depend on risk preferences. In other words, subjects may exhibit risk aversion regarding their subjective distribution about the optimal action. Irrespective of the specific assumptions made, in the beta-binomial theoretical setup, the approximation error from assuming the subject playing the subjective mean is likely small. 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.

signals. Given $\mathbb{E}[S] = a^*$, we define:

$$a^e := \mathbb{E}[a^o] = \lambda \cdot a^* + (1-\lambda) \cdot d. \tag{12}$$

From Section 2, we have $a^* = u^{-1}(\delta^{\Delta t})$. We assume that $u(c) = c^\alpha, \alpha > 0$. To simplify, we may allow δ to absorb α and in effect take $\alpha = 1$ in the proofs. We use the definitions of *RRR* and δ_H introduced in Section 5.2.2. 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:

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

Long horizons, LH, are similarly defined by:

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

Lastly, for our convenience we will at times denote $\delta^{\Delta t}$ by $e^{-\beta \Delta t}$ with $\beta = -\ln|\delta|$.

We now turn to the theoretical predictions underlying the pre-registered predictions spelled out in Section 2.

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. □

The following trivial corollary 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 perfect cognitive precision, $\lambda = 1$, per-period impatience is constant in the time delay. Formally,

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

(ii) For those with imperfect cognitive precision, $\lambda = c + be^{-k\Delta t}$, there exists an interval $[0, T]$ such that per-period impatience decreases in the time delay.

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

for $\Delta t \in [0, T]$ where

$$T = \frac{1}{\beta}\ln\left|\frac{1}{d}\left(1 + \frac{\beta}{k}\right)\right|.$$

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 $\lambda = c + be^{-k\Delta t}$, first note that:

$$\frac{\partial a^e}{\partial \Delta t} < 0$$

when

$$\Delta t < \frac{1}{\beta} \ln \left| \frac{1}{d}\left(1 + \frac{\beta}{k}\right) \right| \tag{25}$$

and that

$$\frac{\partial^2 a^e}{\partial \Delta t^2} > 0$$

when

$$\Delta t < \frac{1}{\beta} \ln \left(\frac{1}{d} + \frac{\beta^2}{k^2 d} + \frac{2b}{kd} \right). \tag{26}$$

Note that (25) is a tighter bound than (26). Hence, when (25) holds we see that:

$$\frac{\partial RRR}{\partial \Delta t} = -\frac{1}{a^e}\frac{\partial a^e}{\partial \Delta t} > 0 \tag{27}$$

$$\frac{\partial^2 RRR}{(\partial \Delta t)^2} = -\frac{1}{a^e}\frac{\partial^2 a^e}{\partial \Delta t^2} + \frac{1}{(a^e)^2}\frac{\partial a^e}{\partial \Delta t} < 0 \tag{28}$$

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

$$\frac{\partial r}{\partial \Delta t} = \frac{\frac{\partial RRR}{\partial \Delta t}\Delta t - RRR}{\Delta t^2} \tag{29}$$

A sufficient condition for (29) to have negative sign is therefore:

$$\Delta t \cdot \frac{\partial RRR}{\partial \Delta t} < RRR \tag{30}$$

We may now define the function:

$$g := RRR - \frac{\partial RRR}{\partial \Delta t}\Delta t \tag{31}$$

and differentiate to find:

$$\frac{\partial g}{\partial \Delta t} = -\frac{\partial^2 RRR}{(\partial \Delta t)^2}\Delta t \geq 0 \tag{32}$$

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

$$g(0) = RRR(0) = -\ln|\lambda + (1-\lambda)d| > 0 \tag{33}$$

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

$$g > 0 \ \Big| \Delta t > 0 \tag{34}$$

substituting in the definition of g shows that (30) is satisfied yielding the result. $\square$

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:*

$$\left.\frac{\partial^2 r}{\partial\lambda\partial\Delta t}\right|_{\lambda=1} > 0 \tag{35}$$

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

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. For $\lambda = e^{-k\Delta t}, c; d \in (0,1)$ we claim:*

$$SA := (RRR_{t_0,t_0+t} + RRR_{t_0+t,t_0+2t}) - RRR_{t_0,t_0+2t} > 0 \tag{36}$$

Since RRR only depends on the time difference and not the start time this is equivalent to considering

$$SA := (RRR_{0,t} + RRR_{t,2t}) - RRR_{0,2t} > 0.$$

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

$$a^e(2t) > a^e(t)a^e(t) \tag{37}$$

Proof. By algebraic manipulations the existence of subadditivity is equivalent to:

$$f(t) := (1-e^{-2kt})d - (1-e^{-kt})^2 d^2 - 2e^{-kt}(1-e^{-kt})\delta^t d > 0. \tag{38}$$

The subadditivity condition is reduced to $f(t) > 0$ for $t > 0$. We first compute that $f(0) = 0$. Accordingly, to prove that $f(t) > 0$ for positive values of t it will suffice to show that $f'(t) > 0$. Taking the derivative yields:

$$f'(t) = 2ke^{-2kt}d(1-\delta^t) + 2ke^{-kt}(1-e^{-kt})d\big(\delta^t + d - \ln|\delta|\delta^t\big). \tag{39}$$

Since $\delta < 1$ we see that the above expression is positive thereby proving the claim.

For constant $\lambda \in (0,1)$, algebraic manipulations reduce subadditivity to

$$g(d) := \lambda(\delta^{2t} - 2d\delta^{t}) + d(1-(1-\lambda)d) > 0. \tag{40}$$

Subadditivity is equivalent to $g > 0 \,\big|\, d \in (0,1)$. We now prove this claim.

Since g is quadratic in d with negative second derivative its unique minima on an interval will be found on the boundary points of the interval, them being, $d \in \{0,1\}$. We note that $g(0)$ is trivially positive. Moving on to the next point we compute

$$g(1) = \lambda(1 + \delta^{2t} - 2\delta^{t}) \tag{41}$$

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

$$h(0) = 0 \tag{42}$$

$$\frac{dh}{dt} = 2\lambda \ln|\delta|(\delta^{2t} - \delta^{t}) \geq 0 \tag{43}$$

Consequently, $g(1) \geq 0$ and may conclude that $g > 0$ for $d \in (0,1)$.

□

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 $SA > 0$. □

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

Note that the non-existence of front-end delay effects in the absence of cognitive noise is driven by our assumption of exponential discounting. Present bias would deliver front-end delay effects in the absence of cognitive noise.

Proof. As mentioned earlier, with exponential discounting, *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 $l > 0$,

$$\Delta FE := RRR_{0,t_2} - RRR_{l,t_2+l} = \ln\left(\frac{\lambda u^{-1}(\delta^{t_2}) + (1-\lambda)d}{\lambda u^{-1}(\delta^{t_2}) + (1-\lambda)d}\right) = 0. \tag{44}$$

□

The following corollary underlies Prediction 4 in the main text:

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

B Additional Figures

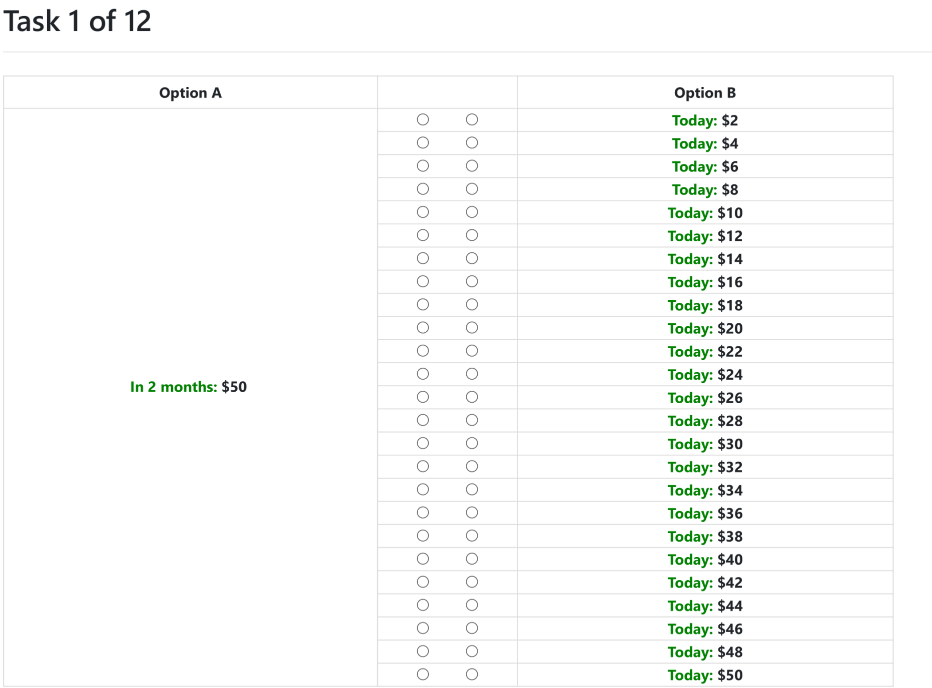

Figure 7: Screenshot of an example decision screen in *Money Main*

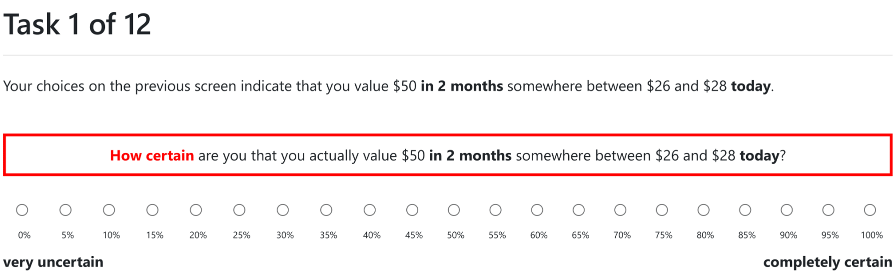

Figure 8: Screenshot of an example cognitive uncertainty elicitation screen in *Money Main*

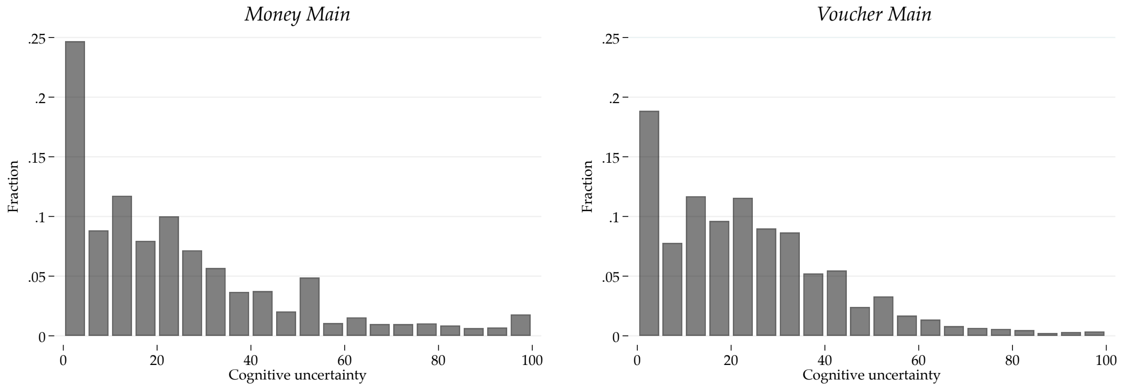

Figure 9: Histogram of cognitive uncertainty statements in *Money Main* (left panel, $N = 7,740$) and *Voucher Main* (right panel, $N = 6,000$).

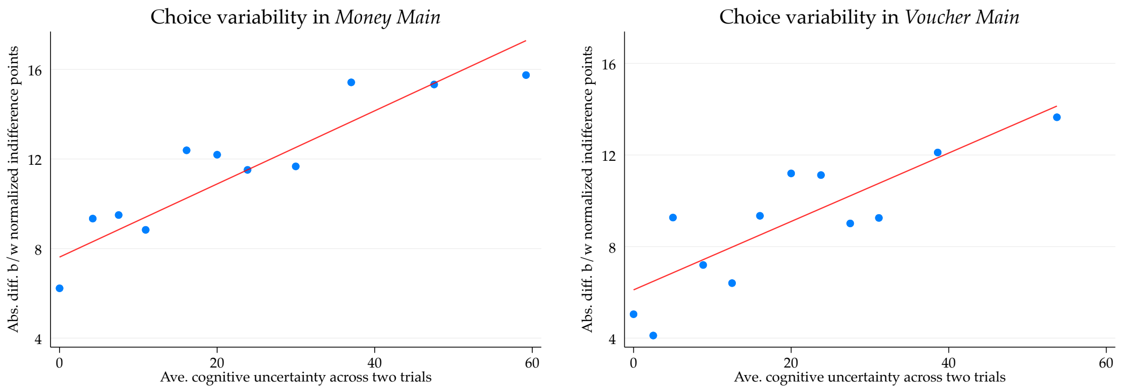

Figure 10: Link between cognitive uncertainty and across-task variability in normalized switch points in an exact repetition of the same decision problem in *Money Main* (left panel, $N = 1,290$) and *Voucher Main* (right panel, $N = 1,000$). The y-axis captures the absolute difference between the normalized indifference points across the two implementations. Average cognitive uncertainty is winsorized at 60 (roughly the 95th percentile in both datasets) for ease of visibility.

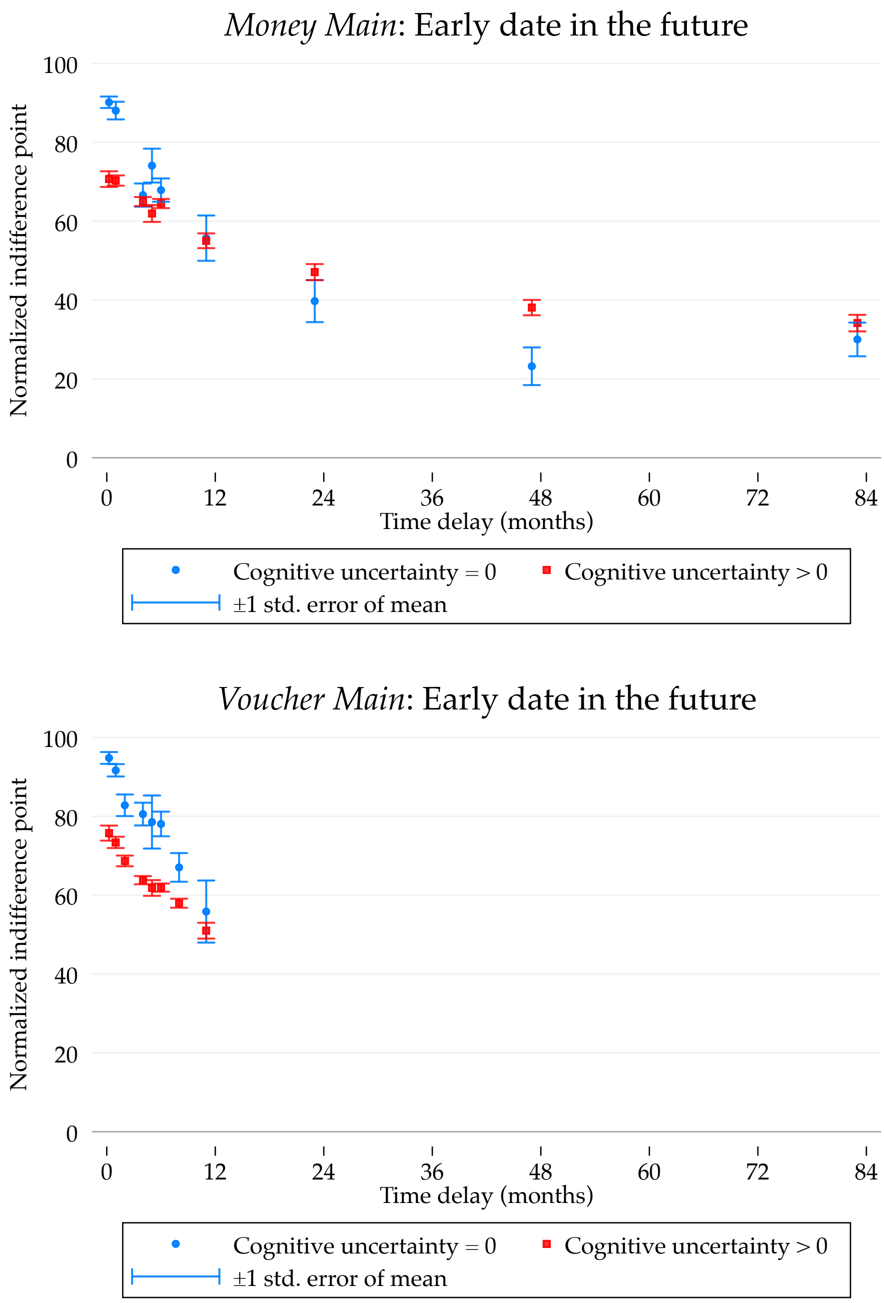

Figure 11: Observed discounting with $t_1 > 0$ in *Money Main* (top panel, $N = 2792$) and *Voucher Main*, $N = 2154$ (bottom panel). The figure shows averages across decisions. Whiskers show standard error bars, computed based on clustering at the subject level.

C Additional Tables

Table 6: 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 3–5
Voucher Main	Incentivized UberEats Voucher Experiments	500	`https://aspredicted.org/b4pw2.pdf`	Sections 3–5
Money Complex Dates	Like *Voucher Main*, except with later payment date displayed as math expression	149	`https://aspredicted.org/77xp6.pdf`	Section 6
Money Complex Amounts	Like *Voucher Main*, except with payoff amount displayed as math expression	153	`https://aspredicted.org/77xp6.pdf`	Section 6
Money Load	Like *Voucher Main*, except with cognitive load manipulation	154	`https://aspredicted.org/77xp6.pdf`	Section 5.2.5
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 5.2.5
Voucher Advice	Like *Voucher Main*, except introducing piece of advice	153	`https://aspredicted.org/jk5s5.pdf`	Section 8

Notes. List of all treatments included in this paper.

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: 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: 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. (9) 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.

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 $t = t_2$, see Figure 12 for an example screenshot.[31] After participants had indicated their indifference amount, the next screen again elicited cognitive uncertainty, see Figure 13.

We here replicate all of our main analyses using these direct elicitation data.

Second, Table 12 documents that cognitive uncertainty is strongly and significantly correlated with impatience over a horizon of one week. Third, columns (1)–(2) of Table 13 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.

Next, Table 14 documents that subadditivity effects strongly increase in cognitive

[31]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.

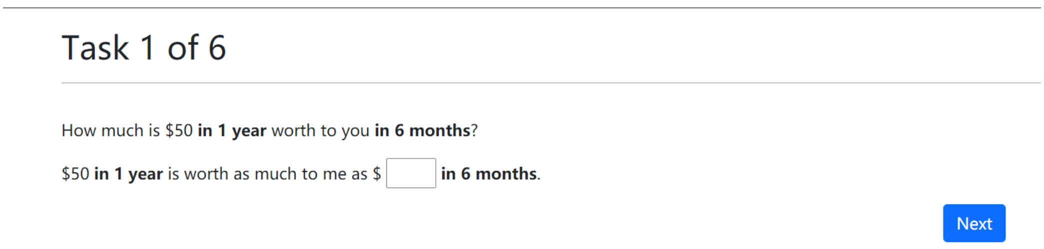

Figure 12: Screenshot of an example decision screen in the direct elicitation part of *Money Main*

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 15 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.

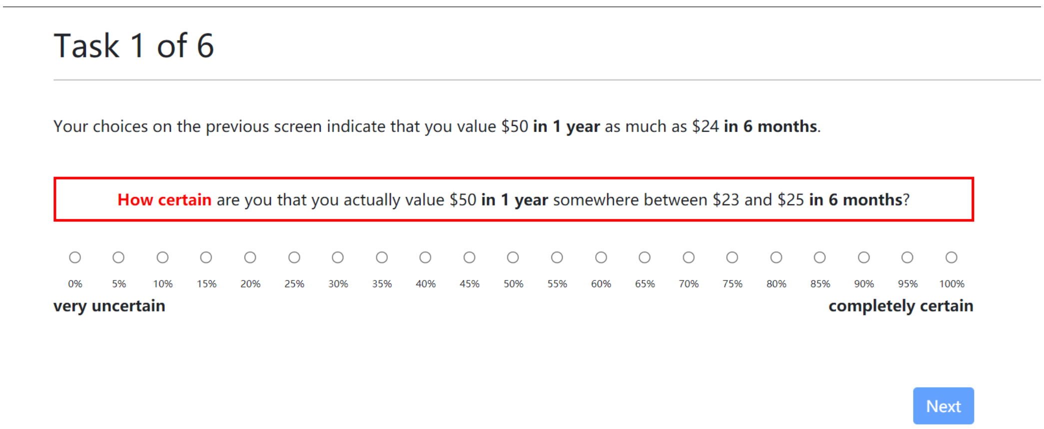

Figure 13: Screenshot of an example cognitive uncertainty elicitation screen in the direct elicitation part of *Money Main*

Table 12: 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 13: 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: 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 15: 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 Experimental Instructions

E.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
	◉	○	**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**

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.

E.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 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**

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.