Optimal Utility: Endogenizing the Cardinal Representation of Riskless Subjective Value in Cognitively Constrained Choosers

Paul W. Glimcher[1], Shreya Sinha[1] and Agnieszka Tymula[2]

August 2025

We investigate the 'optimal representation' of cardinal utility in a cognitively-constrained chooser whose objective is to maximize expected earnings in a riskless setting. We show that optimal utility depends on the *a priori* reward distribution and the level of noise in the nervous system's value encoding process. We quantify the monetary gains and biological costs of relaxing the resource constraint to calculate the optimal level of noise that should be adopted in representation of value, endogenizing both the utility function and the noise term. Our framework links environmental conditions, neural efficiency, and utility. We discuss implications for theory and policy.

JEL codes: D9

Keywords: utility curvature, cognitive costs, neuroeconomics, risk attitude,

---

[1] Grossman School of Medicine, New York University

[2] Corresponding author: Agnieszka.Tymula@sydney.edu.au; School of Economics, University of Sydney

This work was partly supported by the Australian Research Council's Centre of Excellence for Children and Families over the Life Course (Project ID CE200100025). We thank Sura Majeed for excellent research assistance.

Introduction

In the 17th century, Pascal famously argued that people *should* maximize expected value when making decisions and that failing to do so is a mistake – a strategy still employed by nearly all corporate decision-makers. Such an approach, classically, is equivalent to relying on a linear utility function. A key appeal of the linear utility function is its objective accuracy and the long-term maximization of accumulated wealth it provides. It achieves this by preserving objective consistency in pair-wise offer comparisons—for example, the utility difference between $10 and $20 is equal to that between $20 and $30. Despite the simplicity of Pascal's argument, empirical evidence is widely acknowledged to show that people frequently deviate from expected value maximization even for repeated decisions where alternative decisions would lead reliably to higher average earnings. The prevailing explanation for this deviation dates to the 18th century. Bernoulli (1954) explained risk aversion in casino players by arguing that people do not aim to maximize the expected monetary value but rather a concave (logarithmic) expected utility, a function shaped by exogenously determined risk preferences (or the rate of the diminishing marginal utility). Generations of economists and psychologists have built on this foundation with richer notions of what this exogenously determined utility function might look like. Prospect Theory (Kahneman and Tversky 1979), the most often applied form of this approach, describes choosers as having an S-shaped utility function.

One striking feature of the studies of utility conducted over the past three centuries is that economic theorists have generally remained agnostic about the function's origin (exceptions are discussed below and see also a comprehensive review by Vieider (2025)). The standard economic formulation: *choosers are endowed with a stable monotonic utility function,* belies the question of just who does this 'endowing' and for what purpose. It is still a standard in economics that individual preferences are taken as the primitive element of the choice model—an intrinsic feature of the decision-maker, akin to personality traits. This "black box" approach to where the utility function comes from has also defined much of neoclassical economic policy. It suggests that any intervention in decision-making—whether through regulation or behavioral nudges—must be scrutinized, as it may steer individuals away from choices that reflect their own best interests as defined by their inferred idiosyncratic utility functions.

In this paper, taking advantage of the 21st century developments in neuroscience, we (an economist, a neuroscientist, and a computer scientist) highlight what we believe to be a missing piece in the theory – the absence of a specification of just what it is that a reasonable decision-maker *should* be maximizing when facing small-stakes repeated decision problems. The central feature of our approach builds on the recently developed observation that the biological encoding of the subjective value, used by a chooser to guide choice and select an option with a higher value, is constrained by well-documented limitations of the nervous system imposed by the laws of physics and the biological constraints of evolution.[1] Starting from this point we introduce a new, biological definition of cognitive constraints, grounded in neuroscience and random utility theory (McFadden 1974). In addition to assuming that utility is subject to some randomness in the form of an error term, we assume that it is bounded, which is another feature of the neurobiological representation imposed by the laws of physics and observed again and again in neurobiology. We share the core of this biologically accurate definition of limited cognitive capacity with Netzer et al. (2025) but differ from previous approaches that modelled cognitive constraints either through limited number of utility thresholds (Robson 2001) or did not assume a bound on a utility function at all and instead studied the consequences of the limited capacity in decoding noisy utility signals (Woodford, 2012). Like Netzer (2009), we explicitly state that the goal of a chooser is always to go home with more rather than less.[2] With these assumptions, we start off by examining what utility functions maximize total expected earnings in various choice environments as a function of cognitive limitations. Unlike Netzer (2009), who solve for optimal utility in the limit as cognitive capacity approaches infinity, we focus on the curvature of optimal utility for choosers with different levels of cognitive capacity constraint. Formally, we endogenize the utility function as an object employed by the noisy chooser to maximize a well-stated objective function.

Our paper thus contributes to a still relatively small literature that seeks to explain the utility function as serving some optimal goal, such as maximizing fitness instead of treating it as an exogenously endowed enigmatic function (Bucher & Brandenburger, 2022; Frydman & Jin, 2021a; Heng et al., 2020; Netzer, 2009; Netzer et al., 2025; Rayo & Becker, 2007; Robson et al., 2023; Steverson et al., 2019; Stewart et al., 2015)[3]. Independent of differences in theoretical

[1] In the paper, for simplicity, we refer to this "measurement" subjective value function as *utility*, recognizing that this empirically observed cardinal *utility* may differ from other notions of utility in economic theory.
[2] In Glimcher, Sinha, and Tymula (2025), we explicitly compare optimal utilities for choosers whose goal is to maximize earnings and minimize the number of errors.
[3] Another strand of literature (Polanía *et al.* 2019, Heng *et al.* 2020, Frydman and Jin 2021, Khaw *et al.* 2021) focused on the optimal decoding of noisy but unbounded utility signals. See (Vieider 2025) for a comprehensive review.

assumptions, this literature converges on the notion that an optimal utility function should be steeper in regions where choices are made more frequently, and decision mistakes have a greater impact on fitness. We replicate this general finding. We demonstrate that the goal of expected earnings maximization cannot be achieved using a perfectly linear utility function under binding cognitive constraints. Our approach illustrates that curvature in the utility function can be a optimal for maximizing earnings in expectation, as long as the chooser has limited cognitive capacity. Further, we demonstrate that the curvature of the optimal utility function is not fixed but instead must be flexibly determined by the distribution of prizes in the environment in which the chooser operates. These key observations are robust to different modelling approaches and are consistently found across the literature (Vieider, 2025).

We extend the literature by making a series of novel contributions. The closest paper to ours is Netzer et al. (2025). Like Netzer et al. (2025), we assume a bounded utility with a normally distributed error term. Unlike us, in their framework, the curvature in utility is a result of a naïve (imperfect) decoding of a noisy signal, rather than the cognitive limitations in utility encoding. Netzer et al. (2025) solve for optimal utility assuming that cognitive capacity approaches infinity (so it is not a binding constraint) and show that for a naïve decoder, the utility will be S-shaped. Unlike all previous papers, we focus on the biological and cognitive constraint in utility encoding and explicitly explore how the exact level of cognitive capacity affects the selection of an optimal (earnings maximizing) utility function. This extension allows us, for example, to offer a new explanation for the age-related changes in decision errors and utility curvature that are well-documented in the empirical literature.

In addition, we present several novel results by incrementally studying more and more complex choice environments. First, we explicitly manipulate the environmental distribution of prizes from which choice options are drawn (while keeping the choice set size fixed) by independently changing the number of possible prizes in the environment and their probability of entering a choice set. One of the insights that we gain in doing so is that adding more *possible* prizes to the choice environment substantially affects the optimal utility function even when the general distributional pattern of the prize environment (such as uniform, normal, or skewed) is maintained. Our novel insight is that as the number of possible prizes increases, an S-shaped utility function is more likely to emerge as optimal. In line with previous literature, we find that S-shaped utility is optimal for normal-like distributions. This is not surprising because the normal distribution of prizes has been used as the exemplary justification for S-shaped utility at least since Friedman (1989). What is noteworthy is that as the set of possible prizes increases, S-shaped utility can emerge as optimal also for environments with negatively and positively skewed prizes for low cognitive capacity choosers. Second, we manipulate the choice set size separately from the distribution of prizes. We show that increasing the choice set size has an independent effect on the optimal shape of the utility function and that bigger choice set sizes often result in S-shaped utility. This richness of insight is feasible, in part, because in addition to a closed-form solution, our methodology extends to analytical and simulation/numerical approaches (for which we provide the codes for replication and reuse).

Finaly, the major innovation in our paper is that we endogenize cognitive capacity. This is in stark contrast to previous papers that assume that cognitive capacity (if not the utility function) is an exogenously given limitation. We are the first to quantify the monetary gains from increasing cognitive capacity explicitly and we show that these returns are surprisingly low (given that endogenous selection of the earnings-maximizing utility function is allowed) and that the magnitude of these returns diminishes in capacity. We then compare these gains with the costs of additional cognitive capacity to derive the optimal cognitive capacity and utility with no binding constraints on the chooser other than the knowledge of the choice environment. Once one accounts for the cost to increasing cognitive capacity (as required by the laws of physics; see Steverson et al. (2019) for explanation), we show it is optimal to allow for some degree of choice error and utility curvature to maximize monetary earnings in expectation. This implies that even if the costs of reducing noise in the utility function were quite low, what are usually described as *irrationalities* in decision-making might well be expected to persist as optimal despite hundreds of millions of years of evolutionary pressure on the choice mechanism. This may explain why behavioral economists and psychologists have documented so many systematic choice irrationalities in a range of choice environments and in a range of animal species. We note that while modelling the costs of the factors of production is the bread and butter of economic producer theory, this approach has been underutilized in choice theory (but see Friedman (1989) for an early argument for incorporating such analysis to choice theory and (Steverson *et al.* 2019, Bucher and Brandenburger 2022) for a recent examples).

Our approach has important implications for revealed preference theory and what we can infer from observed choices. If observed choices are shaped by an interaction between the choice environment and the cognitive constraints of the chooser, rather than by stable, exogenously given preferences, then the foundation of many economic models—and the policies based upon them—may require reassessment. We propose that rather than "revealed preference theory," a more accurate description of a rational theory of human decision making may be "revealed cognition and decision-environment theory". If our theory is correct, then one could predict the utility (and hence choice structure) of any chooser for any given environment fully known by the chooser.

Our findings contribute to a large body of literature on heterogeneity in observed utility functions which to date largely consists of systematic empirical findings that certain groups make more or less risk averse decisions or show different levels of riskless utility curvature (as measured without risk in random utility theories) than others. Our approach illustrates how environmental factors can lead to concave and convex utilities in the same choosers. We also make predictions about the impact of declining cognitive capacity such as with aging or disease. Our findings align well with the vast literature on the factors that correlate with risk attitudes or utility curvature, such as gender, age, natural disasters, economic shocks, socioeconomic differences, sleep deprivation, and IQ. We extend that literature by making new predictions about how environments affect utility curvature and our framework enables the study of further questions such as whether the costs of cognitive capacity increase with age. We elaborate on how our framework that rests only on the cognitive constraints and distributional properties of the reward environment provides a new understanding of these empirically well-researched heterogeneities in the utility functions in section 5.

Our findings highlight the need for new axiomatic foundations of economics in which decisions (and utilities) adjust to the choice environment. For example, in our framework axioms such as independence of irrelevant alternatives or regularity will often appear to be violated when a choosers' environment changes. Our framework provides the basis for new axiomatic foundations by making precise predictions about the probability of making errors and preference reversals across different choice environments. For example, we predict that the chooser's utility is sensitive to prize distributions, and that the same reward will yield higher utility in environments with positively than negatively skewed prizes, consistent with recent empirical findings (Khaw *et al.* 2017, Frydman and Jin 2021, Guo and Tymula 2021). Additionally, we predict that both an increase in the choice set size and an increase in the number of prizes in the environment independently increase the probability of making an error, a phenomenon known as choice overload. We can also explain why people who live in poorer environments make better decisions when shopping for groceries but will have more difficulty when deciding between high-value financial investments.

Beyond theory, our findings have practical policy applications. If it is true that people's true objective is to maximize expected earnings and the observed choices are not a reflection of exogenously given preferences but instead are due to a combination of cognitive limitations and environment, then economists may reconsider their approach to policy. We discuss some possible avenues in section 4.

Following McFadden's classical approach (McFadden 1974), we define utility curvature using choice stochasticity, meaning that our analysis is conducted entirely in the domain of riskless choice. While our findings appear to have clear implications for risky choice, it should be noted that we do not address lotteries as choice objects directly in the results section. In the discussion, we explore an extension of this work to risky choice.

The paper proceeds as follows: Section 2 presents our theoretical framework, outlining how cognitive constraints shape utility functions. Section 3 develops our main results, demonstrating the impact of different environmental factors and cognitive constraints on optimal decision-making. Section 4 puts our theoretical work in the context of well-known empirical findings on the heterogeneity in utility and discusses policy implications. Finally, in section 5 we conclude by making further connections to existing research, and potential future directions.

2. Theoretical Framework

Consider a general framework with a chooser whose objective is to maximize total earnings over all decisions. For simplicity, we assume that the outcomes of all decisions are immediate allowing us to abstract away from discounting. The set of possible riskless prizes is defined by their objective values as $X = \{x_0, x_1, \dots x_n\}$ and all prizes are equally spaced in terms of their objective value such that $\forall i > 0,\ x_i - x_{i-1} = a$ where $a \in \mathbb{R}_+$. Choice sets are subsets of X with the prizes in each choice set randomly and independently drawn from a specified probability distribution. This distribution thus defines which prizes are abundant, which prizes the chooser is more likely to encounter in choice sets, and which prizes are expected to occur rarely. The chooser picks one alternative from a choice set by maximizing a random and subjective utility function $U(x) = u(x) + \epsilon$, with a bound on $u(x) \in [0, k]$ ($k \in \mathbb{R}_+$) and the noise term, ϵ, which follows a standard normal distribution, $N \sim (0, \sigma^2)$. [4] We impose the restriction that the utility function cannot be infinite. This comes from the constraint that utility must be instantiated in the brain which is a system constrained by the laws of physics. This is required of the brain because the number of neurons in the brain is finite and the amount of information that each neuron can carry is finite, imposed by physical law and verified by empirics, as well. We make no assumptions here about how binding that constraint is, just that it exists. As in any random utility model, it is possible that $U(x_j) < U(x_i)$ when $x_j > x_i$ and hence by maximizing U the chooser may pick an option with a lower objective value. We call such instances errors.

Definition. *The chooser has a limited capacity (c) defined as* $c = \frac{k}{\sigma}$.

The capacity limit, *c*, is defined as the ratio of the maximum value that $u(x)$ can take (*k*) and the standard deviation of the noise term (σ). In plain words, it captures how many standard deviations of noise, the chooser can fit within $u(x)$ bounds of 0 and k. The precision with which the chooser distinguishes between different options and hence the probability of making an error is directly related to capacity. Other things being equal, the smaller the capacity, the more likely it is that the chooser makes an error.

Remark: *With a goal to maximize the earnings, the chooser can make fewer errors by taking advantage of the whole* $u(x)$ *range and setting* $u(x_0) = 0$ *and* $u(x_n) = k$.

In the next section, we use a combination of analytical and computational methods to establish the $u(x)$ that allows the chooser to achieve their objective of maximum total earnings.

3. Optimal Utility

3.1. Closed-form Solution for a Simple Case

[4] The cumulative distribution function (CDF) of this distribution is denoted by $\Phi(.)$ and the probability distribution function (PDF) as $\phi(.)$.

Assume there are three prizes $X = \{x_0, x_1, x_2\}$ and choice sets are always binary. The possible choice sets $\{x_0, x_1\}$, $\{x_0,x_2\}$, and $\{x_1,x_2\}$ occur with respective probabilities $p_{01}, p_{02},$ and p_{12}.

Proposition 1: To maximize earnings, the chooser sets $u(x_0) = 0$, $u(x_2) = k$, and

$$u(x_1) = \begin{cases} 0 & if \ \frac{1}{2k}(k^2 + 4\sigma^2 ln\frac{p_{01}}{p_{12}}) < 0 \\ \frac{1}{2k}(k^2 + 4\sigma^2 ln\frac{p_{01}}{p_{12}}) & if \ 0 < \frac{1}{2k}(k^2 + 4\sigma^2 ln\frac{p_{01}}{p_{12}}) < k \\ k & if \ \frac{1}{2k}(k^2 + 4\sigma^2 ln\frac{p_{01}}{p_{12}}) > k \end{cases}$$

Proof

To show: derive the optimal $u(x_1)$. Maximizing earnings is equivalent to minimizing the objective value of expected loss due to errors. In our simple case the chooser can make three types of errors. First, the chooser can choose x_0 from $\{x_0, x_1\}$, an error which happens with probability $P(U(x_0) > U(x_1)) = 1 - \Phi(\frac{u(x_1)-u(x_0)}{\sqrt{2\sigma^2}}) = 1 - \Phi(\frac{u(x_1)}{\sqrt{2\sigma^2}})$. Second, the chooser can choose x_1 from $\{x_1, x_2\}$ which happens with probability $P(U(x_1) > U(x_2)) = 1 - \Phi\left(\frac{u(x_2)-u(x_1)}{\sqrt{2\sigma^2}}\right) = 1 - \Phi\left(\frac{k-u(x_1)}{\sqrt{2\sigma^2}}\right)$. The objective cost of each of these two types of error is equal to $a \equiv x_i - x_{i-1}$. Third, the chooser can choose x_0 from $\{x_0, x_2\}$ with probability $P(U(x_0) > U(x_2)) = 1 - \Phi\left(\frac{u(x_2)-u(x_0)}{\sqrt{2\sigma^2}}\right) = 1 - \Phi\left(\frac{k}{\sqrt{2\sigma^2}}\right)$. Note that the probability of the third type of error is unaffected by $u(x_1)$ and therefore we disregard it.

Let $u(x_1) = y_1$. The chooser selects $u(x_1)$ to minimize the expected cost of making the first two types of errors by solving:

$$\min_{y_1} p_{01}(1 - \Phi\left(\frac{y_1}{\sqrt{2\sigma^2}}\right) a + p_{12}(1 - \Phi\left(\frac{k - y_1}{\sqrt{2\sigma^2}}\right) a$$

The first order condition yields:

$$\frac{\partial}{\partial y_1}(p_{01}(1 - \Phi\left(\frac{y_1}{\sqrt{2\sigma^2}}\right) a + p_{12}(1 - \Phi\left(\frac{k-y_1}{\sqrt{2\sigma^2}}\right) a)=0 \Leftrightarrow$$

$$-a\frac{\partial}{\partial y_1}(p_{01}\Phi\left(\frac{y_1}{\sqrt{2\sigma^2}}\right) + p_{12}\Phi\left(\frac{k-y_1}{\sqrt{2\sigma^2}}\right))=0 \Leftrightarrow$$

$$-a(p_{01}\phi\left(\frac{y_1}{\sqrt{2\sigma^2}}\right)\frac{\partial}{\partial y_1}\left(\frac{y_1}{\sqrt{2\sigma^2}}\right) + p_{12}\phi\left(\frac{k-y_1}{\sqrt{2\sigma^2}}\right)\frac{\partial}{\partial y_1}\left(\frac{k-y_1}{\sqrt{2\sigma^2}}\right))=0 \Leftrightarrow$$

$$-a(p_{01}\frac{1}{\sqrt{2\pi}}e^{-\frac{y_1^2}{4\sigma^2}}\frac{1}{\sqrt{2\sigma^2}}-p_{12}\frac{1}{\sqrt{2\pi}}e^{-\frac{(k-y_1)^2}{4\sigma^2}}\frac{1}{\sqrt{2\sigma^2}})=0 \Leftrightarrow$$

$$p_{12}e^{-\frac{(k-y_1)^2}{4\sigma^2}}=p_{01}e^{-\frac{y_1^2}{4\sigma^2}} \Leftrightarrow$$

$$e^{-\frac{(k-y_1)^2}{4\sigma^2}+\frac{y_1^2}{4\sigma^2}}=\frac{p_{01}}{p_{12}} \Leftrightarrow$$

$$e^{\frac{2ky_1-k^2}{4\sigma^2}}=\frac{p_{01}}{p_{12}} \Leftrightarrow$$

$$\frac{2ky_1-k^2}{4\sigma^2}=ln\frac{p_{01}}{p_{12}} \Leftrightarrow$$

$$y_1=u(x_1)=\frac{1}{2k}(k^2+4\sigma^2 ln\frac{p_{01}}{p_{12}})$$

■

Although our setting is simple, Proposition 1 provides the first set of key insights. First, it demonstrates how the probability distribution of prizes affects the optimal $u(x)$ in a capacity-constrained chooser. The larger the probability of encountering $\{x_0, x_1\}$, relative to $\{x_1, x_2\}$, the larger value of the optimal $u(x_1)$. This occurs because it is optimal to increase the distance between the utilities and thus decrease the number of erroneous choices in the more frequently encountered choice set. Note that when choice sets $\{x_0, x_1\}$ and $\{x_1, x_2\}$ are equally likely ($p_{01}=p_{12}$), in this simple case, the optimal solution simplifies to $u(x_1)=\frac{k}{2}$ meaning that optimal $u(x_1)$ is exactly in the middle of the range of possible values of $u(x)$. Second, we note that low capacity can push $u(x_1)$ to the boundary when $p_{01}\neq p_{12}$. For example, when $\frac{p_{01}}{p_{12}}>1$, the lower the upper bound k and/or the higher the standard deviation of noise (σ), the more likely it is that $u(x_1)=k$ (because it is more likely that $\frac{1}{2k}(k^2+4\sigma^2 ln\frac{p_{01}}{p_{12}})>k$). When $\frac{p_{01}}{p_{12}}<1$, the lower the bound k and the higher the standard deviation of noise (σ), the more likely it is that $u(x_1)=0$, because it is more likely that $\frac{1}{2k}(k^2+4\sigma^2 ln\frac{p_{01}}{p_{12}})<0$. This means that more cognitively constrained choosers are more likely to make fully random decisions for some values of prizes. Third, as the capacity increases, in the limit the optimal $u(x)$ approaches a linear function, even when $p_{01}\neq p_{12}$. Trivially, if $u(x)$ is noiseless, any $u(x)$ that preserves the preference ordering of prizes is equally optimal. Finally, notice that the optimal solution does not depend on the cost of a mistake a. This implies that in this specific simple environment, maximizing earnings and minimizing the number of errors yields the same optimal $u(x)$. This implies that a researcher who is interested in figuring out whether choosers maximize earnings or minimize errors, will not be able to tell the difference in environments with three prizes and binary choice sets. In an accompanying psychology paper. Glimcher, Sinha, and Tymula (2025), we show that this feature does not generalize to more complex environments in which it is possible to make such a distinction, as suggested by some previous work (Heng *et al.* 2020).

The key insights from Proposition 1 are illustrated in Figure 1. The three panels A-C illustrate choice environments characterized by different distributions of prizes (top panel) and the resulting optimal $u(x)'s$ that maximize the earnings in these environments for different capacity levels (bottom panel). Each of the top figures illustrates the probability with which each of the three prizes enters a binary choice set. In panel A, the prizes are distributed such that prizes with the lower objective value are more likely to enter a choice set. In panel B, each prize is equally likely, and in panel C, the prizes with higher objective value are more likely to be in the choice sets. The bottom panel illustrates the optimal $u(x)$ for each environment at different capacity levels.[5] What is striking is that even though the chooser has the same objective in these three environments (which is to maximize their average earnings), the optimal $u(x)$'s are different, particularly at low capacities (darker curves). The environment with positively skewed prizes, generally produces concave optimal $u(x)$. The environment with uniform distribution of these three prizes, produces a linear $u(x)$. Finally, the environment with negatively skewed prizes, produces convex $u(x)$. As capacity increases, in each environment, the optimal $u(x)$ becomes more linear and thus closer to standard expected value maximization which assumes unlimited capacity. In the following sections, we will examine which of these predictions hold as we extend to more complex environments.

[5] To draw these figures, we assumed $\sigma = 1$, so that $c = k$.

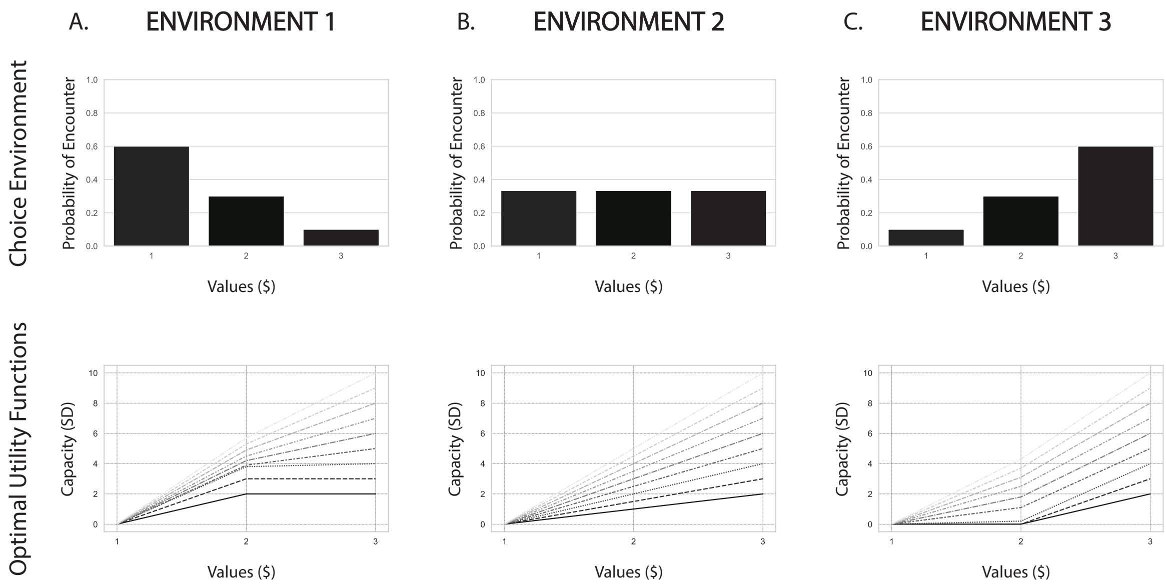

Figure 1. Closed-Form Solution Optimal Utilities

3.2 Analytical Solution - Any Number of Prizes and Binary Choice

For simplicity in the remainder of the paper, without loss of generality, we will assume $\sigma = 1$ which means that capacity and upper bound of $u(x)$ are equal, $c = k$. This allows us to draw optimal $u(x)$ in units of capacity without introducing additional notation. We do this because we prefer this unit of utility as it comprehensively captures the effects of both k and σ on precision and probability of making errors. We assume $\sigma = 1$ to avoid the confusion that could arise from the introduction of unnecessary notation for a utility with a different cardinal unit. This means that the cardinal units of utility for the rest of the presentation are in standard deviations, a representation originally captured by standard random utility theory (McFadden 1974). However, readers should keep in mind that an increase in capacity can be a result of an increase in k, decrease in σ, or both. This assumption of course implies that to take advantage of the whole available range of $u(x)$, the chooser will set the utility of the highest value prize as $u(x_n) = c = \frac{k}{\sigma}$.

Let us now consider a slightly more complicated choice environment in which the set of possible prizes is larger than 3 ($|X| > 3$) and the choice sets are still binary. Given that there are now more error types (because there are more possible binary choice sets), we solve for optimal $u(x)$ using a different procedure. Here, without loss of generality, we assume that for all i, the objective value of $x_i = i$ and $a = 1$.

Our procedure, coded in Python, is as follows: First, we define the environment by a finite set of possible prizes X and a probability distribution P that determines how likely each of the prizes is to enter a binary choice set. Second, we assume a capacity limit c. Third, we set $u(x_0) = 0$ and $u(x_n) = c$. Our task is to find optimal $u(x_i)$ for all $i \in (0, n)$. Given that there is an infinite number of candidate optimal utility functions, to make the problem manageable, we restrict the range of possible utility values to be in steps of 0.1 which makes the set of candidate $u(x)$ finite. Next, for each candidate $u(x)$, we analytically calculate expected earnings in the assumed environment and pick the $u(x)$ with the highest observed expected earnings as optimal.

We calculate the expected earnings for each candidate $u(x)$ using the following steps. First, for each possible binary choice set, we calculate the expected objective value (not utility) of decisions made in this choice set using the current candidate $u(x)$ at the assumed capacity level. Because the random utility error term for a given capacity is normally distributed and the choice sets are binary, we can define the probability of making an error when choosing from $\{x_i, x_j\}$ as $1 - \Phi(\frac{u(x_j)-u(x_i)}{\sqrt{2\sigma^2}})$ and compute it directly in Python. We thus calculate the expected earnings (EV) of choice set $\{x_i, x_j\}$ under a given $u(x)$ as $EV_{i,j} = (1 - \Phi\left(\frac{u(x_j)-u(x_i)}{\sqrt{2\sigma^2}}\right))x_i + \Phi\left(\frac{u(x_j)-u(x_i)}{\sqrt{2\sigma^2}}\right)x_j$. Once the expected earnings of decisions in each possible choice set are computed for a given $u(x)$, we add them up weighing them by the probability with which each choice set occurs given the distribution of choice problems for the current environment. We repeat the exercise until we compute total expected earnings under each candidate $u(x)$ and then pick the $u(x)$ with the highest earnings as optimal. To understand how environments and capacity affect optimal $u(x)$, we then repeat the exercise for different environments and for different capacity levels.

In Figure 2, we illustrate optimal $u(x)$ in environments with $|X| = 4$ with positively skewed (Figure 2A), uniform (Figure 2B), and negatively skewed (Figure 2C) prize environments. The key insights are that, in the positively and negatively skewed environments, we see similar patterns as in our closed-form result with just three prizes (illustrated in Figure 1 in the preceding section). In positively skewed environments, where low value prizes are more likely to enter choice sets, concave $u(x)$'s emerge as optimal (Figure 2A). In negatively skewed environments, where high value prizes are more likely to enter choice sets, convex $u(x)$'s emerge as optimal (Figure 2C). The uniform environment is different though. If the capacity is low enough (darker curves), we find that an S-shaped $u(x)$ is optimal.

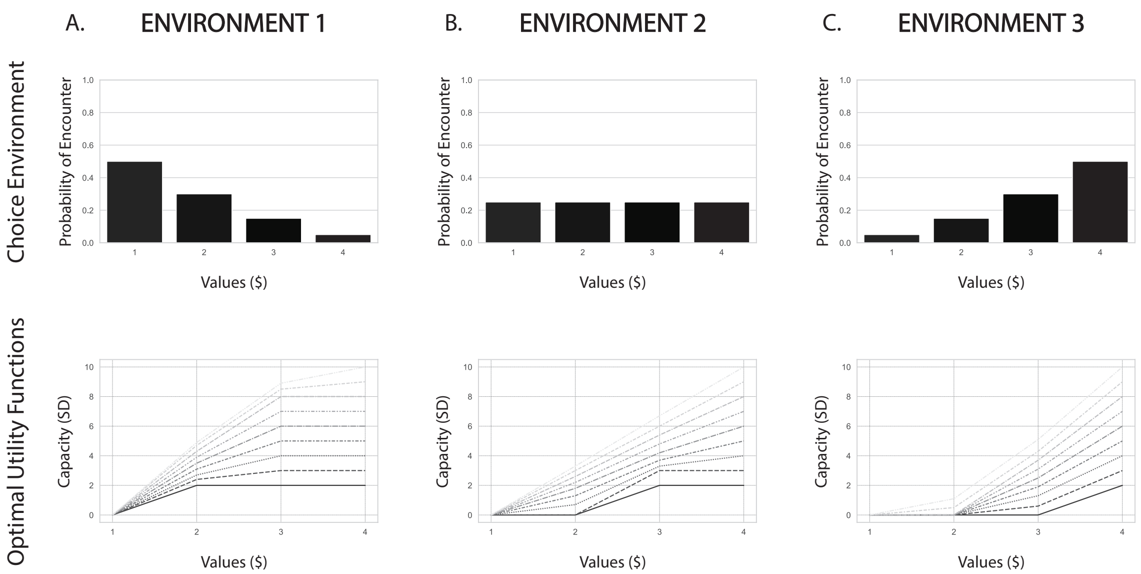

Figure 2 Optimal utility with Four Prizes

Result: *For sufficiently low capacities, environments with positively skewed prizes, result in concave optimal* $u(x)$. *Environments with negatively skewed prizes, result in convex optimal* $u(x)$. *Environments with uniformly distributed prizes, result in S-shaped optimal* $u(x)$.

To gain further insights into how increasing the size of the set of possible prizes X (while holding the choice set fixed as a binary choice problem) impacts optimal $u(x)$, we repeat our numerical procedure for $|X| = 7$. The results are illustrated in Figure 3 and give two key insights. First, with a larger set of possible prizes, for all distributions at sufficiently low capacities, S-shaped $u(x)$'s emerge as optimal. This happens not only in the uniform environment, but also in the positively and negatively skewed environments. In the positively skewed environments, as capacity increases, the optimal $u(x)$ changes from S-shaped to concave and then towards the linear function. In the negatively skewed environment, as capacity increases, the optimal $u(x)$ changes from S-shaped to convex and then linear. The second key insight is that when capacity is low enough and the optimal $u(x)$ is S-shaped, the distribution of prizes in the environment determines the point at which the $u(x)$ inflects. We note that this point would traditionally be called a reference point in behavioral economic theory. Rather than being a fixed number (e.g. status quo), or a rule (e.g. maxmin), our results suggest that the inflection point in the utility curvature, if it occurs, changes according to the distribution of prizes in the environment. As illustrated in Figure 3, the inflection point occurs at higher prize values when higher value prizes are more frequent. *It is important to note though, that in our theoretical framework, the reference point is not a primitive, or even a meaningful element of our approach.* We note this because the reference point is a prominent primitive in many behavioral choice models. In our framework one can determine in what environments and at what capacities this point of maximal inflection in utility curvature emerges as an optimal representational feature.

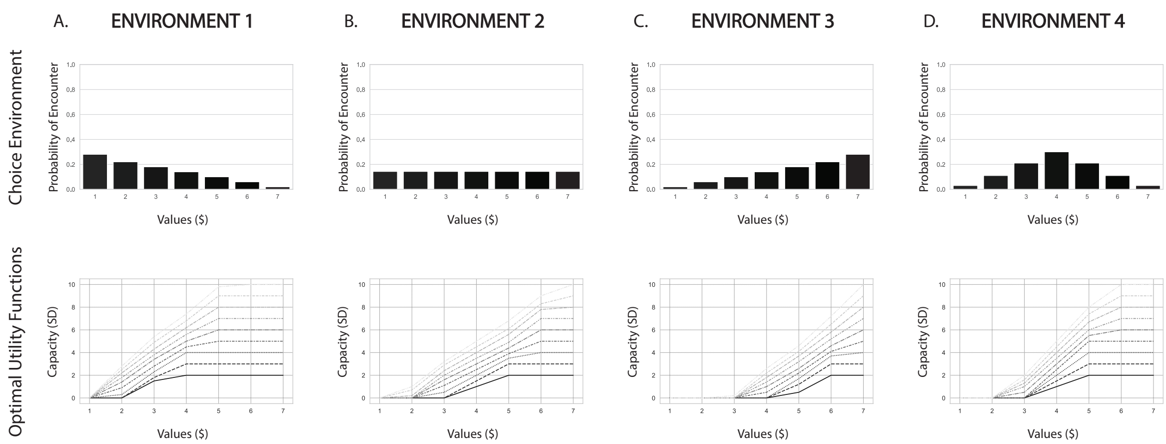

Figure 3 Optimal Utility with Seven Prizes

Result: *As the set of prizes increases, S-shaped $u(x)$ are more common. The inflection point in $u(x)$ depends on the distribution of prizes.*

3.3 Monte Carlo Solution for Any Environment – Larger Choice Sets

In the preceding section, we investigated how increasing the size of the set of prizes in the environment changes the optimal $u(x)$ and established how the optimal $u(x)$ is determined by the underlying distribution of prizes. Now we ask whether, for a given distribution of prizes, increasing the choice set size (the number of options between which the chooser selects in a given choice problem) has any effect on the shape of the optimal $u(x)$. A computational approach with Monte-Carlo simulations allows us to handle the complexity of this problem, ensuring reasonably accurate estimations of expected payoffs by assessing a million simulated choices for each possible combination of utility function, noise, environments and choice set size.

As in the previous procedure, first, we define the environment by a finite set of possible prizes X and a probability distribution that determines the likelihood of each prize entering a choice set. Here, we additionally specify the size of the choice set. Second, we assume a capacity limit c. Third, we set $u(x_0) = 0$ and $u(x_n) = c$. Our task is to find optimal $u(x_i)$ for all $i \in (0, n)$. Given that there are an infinite number of candidate optimal utility functions, to make the problem manageable, we restrict the range of possible utility values by limiting the number of possible $u(x)$ to 50 values, equally spaced from 0 to c which makes the set of the candidate $u(x)$ finite. Next, for each candidate $u(x)$, we calculate its average expected earnings in the assumed environment using a Monte-Carlo simulation. For each $u(x)$, we simulate choices in 1,000,000 randomly drawn choice problems and calculate average earnings. The prizes that form the individual choice problems are drawn from the set of all prizes according to the assumed probability distribution. To determine which prize is chosen from the choice problem presented to the chooser, we add a Gaussian noise $\epsilon \sim N(0,1)$ to $u(x)$ and assume that the simulated chooser picks the option with the highest $U(x)$. The optimal utility function is selected as the one with the maximum average expected earnings.

Figure 4 illustrates that the optimal $u(x)$ changes as the choice set size increases from choosing between 2 to 10 options in a positively skewed environment for a chooser with capacity $c = 4$. The key insight is that as the choice set size increases, the steepness of the optimal $u(x)$ across the prize ranges changes. As the choice set size increases, the $u(x)$ becomes shallower for low value prizes and steeper for high value prizes. This adjustment is optimal, because as choice set size increases, the choice sets are more likely to include high value prizes. From an optimality perspective, if a choice set includes both high value and low value prizes, it is always better for the chooser to be able to precisely discriminate between the high value prizes, that are effectively chosen, rather than low value prizes (that are unlikely to be selected anyway) in a given choice set. In Figure 5, we illustrate this effect of increasing choice set size for different environments and capacities.

Result: *The size of the choice set affects the optimal* $u(x)$*. As choice sets increase, the optimal* $u(x)$ *is steeper for high value prizes and shallower for low value prizes.*

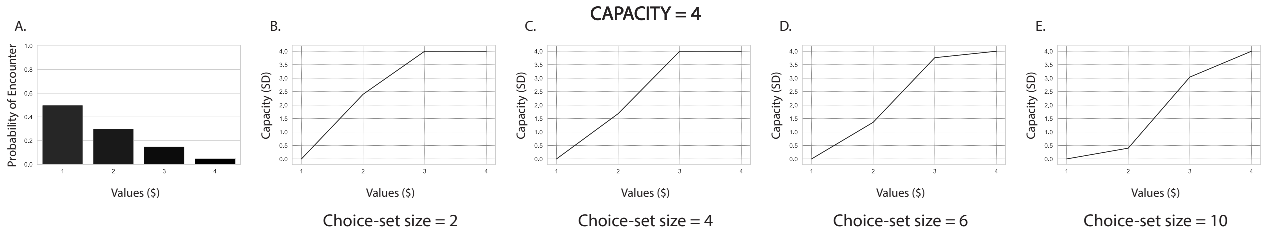

Figure 4 Optimal Utility as Choice Set Size Increases

One interesting feature is that as the number of options presented in a choice problem increases, the conditional probability that a given x-value will be selected changes. This is, in essence, a violation of the regularity axiom of McFadden (1974) that is observed as we move across choice set sizes, a violation driven by optimality. This can be seen clearly in Figures 4 and 5. Increases in the choice set size change the utility of each *x*, affecting the relative probabilities with which the prizes would be selected when offered. For example, using the utilities in Figure 4, consider the relative probability of choosing 3 versus 4 when they are both offered in a choice set. When the choice set size is equal to two, the utility function is horizontal, and therefore the probability of choosing each prize is the same and equal to 50%. However, as the choice set size increases to ten, the utility function is no longer horizontal and the probability of choosing 4 is higher than the probability of choosing 3.

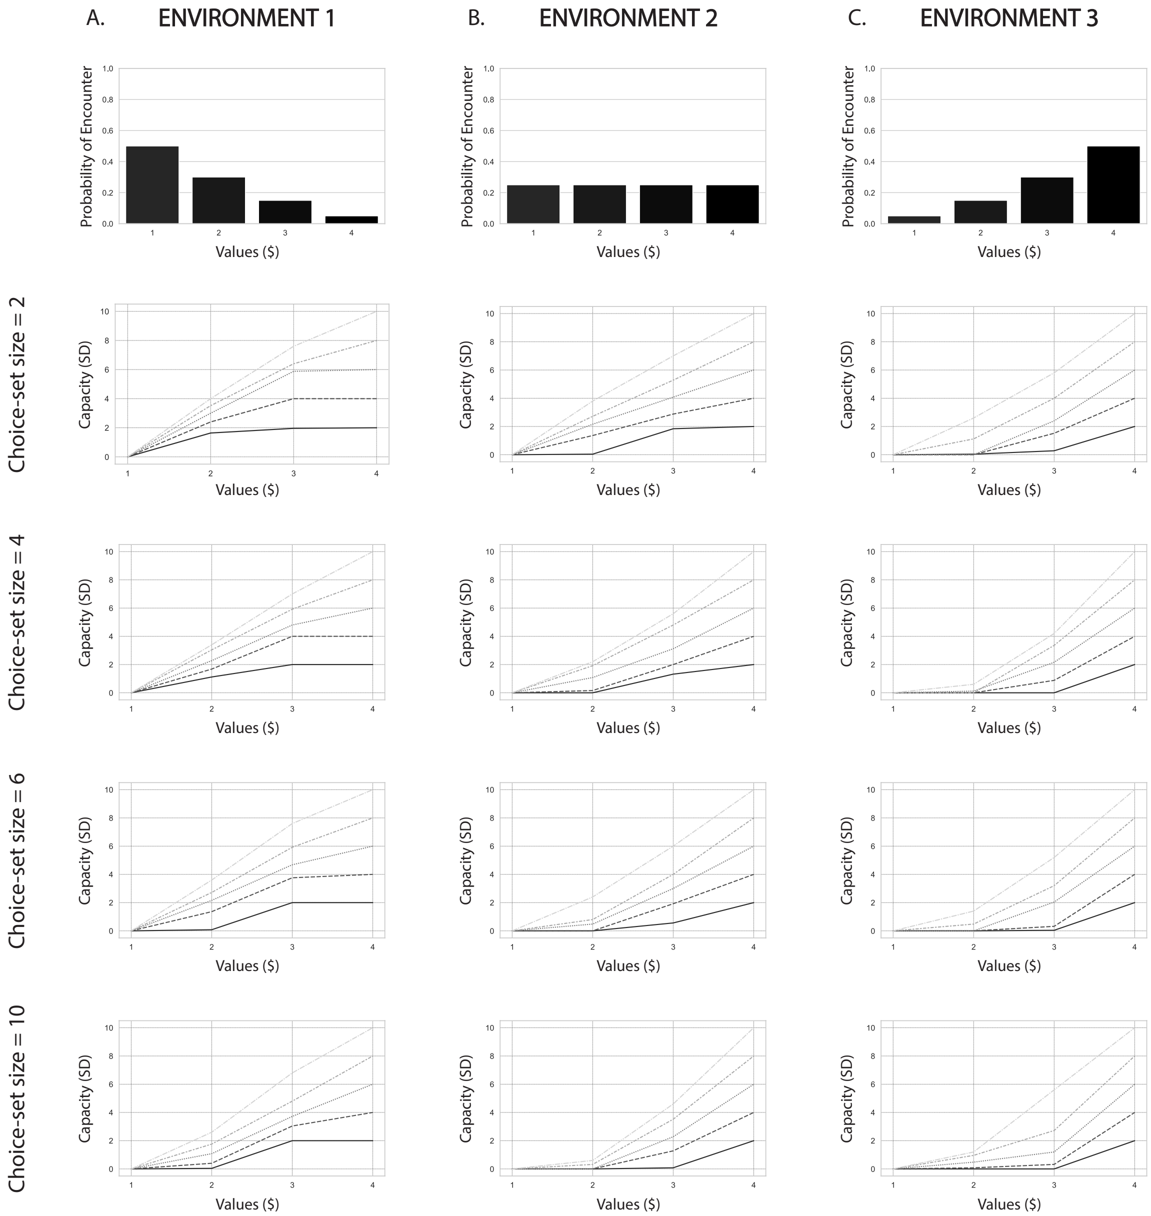

Figure 5 Optimal Utility as Choice Set Size Increases by Capacity

3.4. Deriving Optimal Capacity

So far, we have considered capacity as a largely fixed parameter, examining how a chooser would behave if she was exogenously endowed with a range of possible capacities. We established that *when* the capacity constraint is binding, a chooser whose objective is to make as much money as possible, will nevertheless have curved utility function that depends on the distribution of rewards in their environment, choice set size, and their capacity level. Given the broad empirical evidence that choosers rarely have linear utilities, we take this to imply that at an empirical level the capacity constraint must be binding for real human chooser. This raises an important question: why is capacity limited in the first place, and why have we not evolved to make decisions that consistently maximize expected earnings by developing higher cognitive capacities?

To begin to examine that issue we start by noting that from a biological perspective capacity has been shown to be surprisingly costly (Lennie 2003, Glimcher 2022). Briefly, the capacity of the nervous system is constrained by two features: The number of neurons in the brain and the total informational capacity, or number of *action potentials*, generated by each of those individual neurons. A combination of these two factors defines k in our model. How costly then is a given brain with the maximum k it can produce? How can reducing the number of action potentials reduce costs? Hypothetically, one could envision a brain with so many neurons or neurons that can produce action potentials at such fast rates, that the capacity constraint would not be binding. However, it is essential to acknowledge that this would come at greater metabolic cost. Just as our muscles require calories to produce movement, our neurons consume calories to produce action potentials, and neural activity is particularly expensive. Although the brain accounts for approximately only 3% of our body mass, it consumes 20% of an average chooser's recommended daily caloric intake. Increasing the processing capabilities of our brains ten times by adding more neurons or more action potentials, would push this figure to something like 70% (Glimcher 2022). This naturally reframes optimal capacity as a resource allocation problem: Given that additional capacity is costly and yields limited gains, how much capacity would a chooser elect to allocate to a given problem, and what limits on capacity would evolution place on choosers?

To solve for optimal capacity, we continue to assume that the chooser's goal is to maximize their earnings. For simplicity, we assume that the cost of capacity increases according to a simple linear function $cost(c) = ac$, where $a > 0$ is a cost parameter and $c > 0$ is the capacity level. For illustration, we employ five values of the cost parameter a (from 0.03 to 0.13) to demonstrate the relationship between costs and optimal capacity. These five cost functions are illustrated in the middle panel of Figure 6. We simplify the choice problem to binary choice sets and consider three environments (Figure 6; A: positively skewed, B: uniform, and C: negatively skewed). Using the analytical approach specified in section 3.2, for integer cognitive capacity values that range from 1 to 10, we solve for the optimal $u(x)$ and calculate the corresponding expected earnings per decision.

These earnings that reflect the most the chooser could earn in expectation in each environment, for each capacity, are shown as a function of capacity in the second to top panel of Figure 6. It is important to note that each of the points on the earning curves employs a different utility function, one optimally selected for that capacity level and that environment. While the earnings of the optimal chooser are consistently higher than those of a random chooser (indicated by the horizontal gray line), it is striking that the difference is surprisingly modest and largely plateaus at capacities of 4 or greater. Expected earnings increase with capacity, but at a diminishing rate, reflecting the diminishing marginal returns to increased capacity.

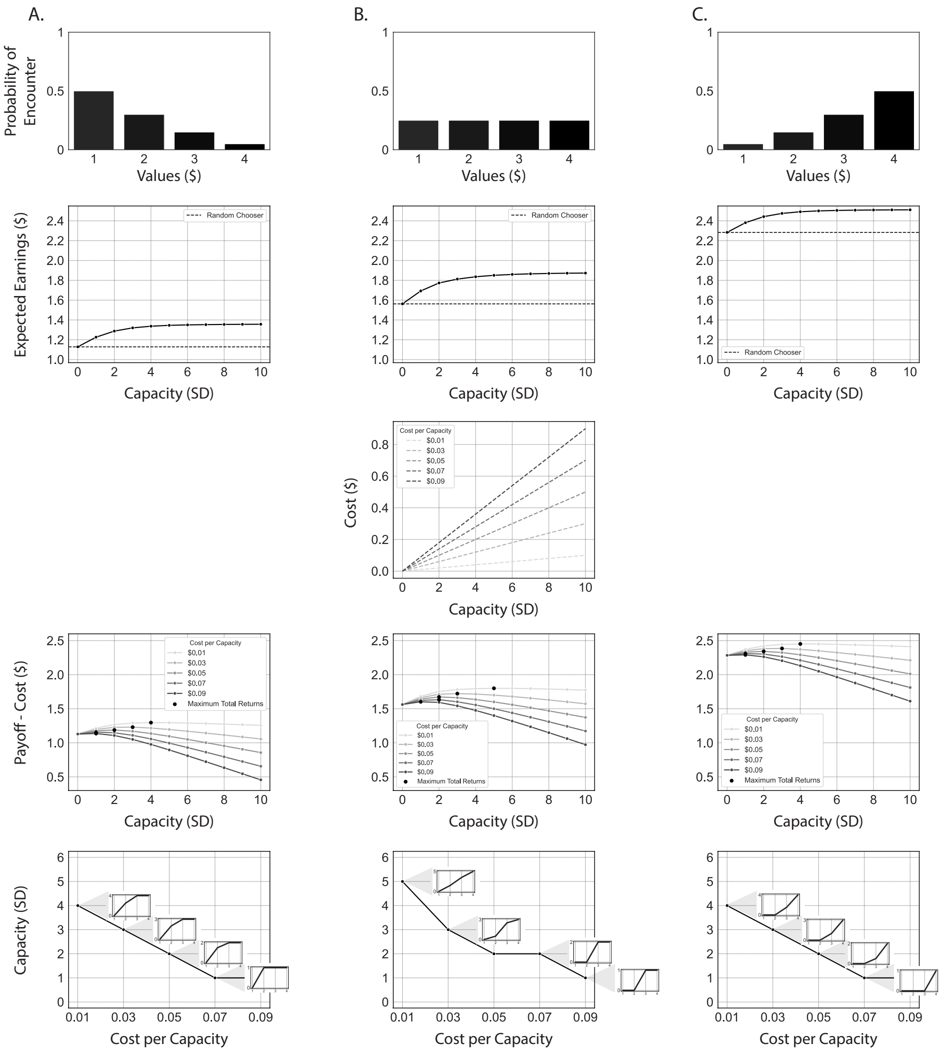

Figure 6 Optimal Capacity for Different Cost Functions

Although the returns to capacity are diminishing, if adding capacity were truly costless, it would of course be optimal for a chooser to have infinite capacity. However, as we noted at the beginning of this section, capacity is by definition costly (Steverson *et al.* 2019). Therefore, to determine the optimal capacity level (and thus optimal $u(x)$), we must consider these costs. We combine the benefits and costs of increasing capacity by calculating the difference between expected earnings and costs under each capacity level. These net earnings are illustrated in the second-to-bottom panel of Figure 6 with darker curves corresponding to highest costs and the capacity that yields highest net earnings marked with a black dot. The black dots thus indicate the optimal capacity levels for each cost function, i.e. the capacity at which the chooser gets the maximum earnings from their decision after accounting for the costs of capacity.

The bottom panel of Figure 6 brings together the full analysis. Here we plot, for each cost level, the optimal capacity (and optimal $u(x)$ as insets). Were one to know the actual cost of capacity, it would be simple to select the optimal capacity and optimal $u(x)$ for any given environment and choice set size. More realistically, we can gain additional insights even without knowing the actual cost of capacity. Were one to observe stable utility functions in a real chooser and the prize distribution in which this chooser operates, one could then infer the cost of capacity. One could imagine strengthening such an inference by systematically increasing either the choice set size or the payoff magnitudes and observing how the estimated utility functions change. Such an approach would, in principle, allow for an estimation of the actual capacity cost faced by the chooser.

4. Empirical Implications and Policy

Unlike in the traditional models that take preferences as primitives, the primitives of our framework are the distribution of the rewards in the environment and the constraint that the cognitive capacity of the chooser is finite. Starting from this point we make a series of predictions of how each affects the shape of the optimal utility function for maximizing expected earnings. Here, we discuss the existing empirical evidence that supports our theoretical findings, highlight where evidence is still missing, and importantly elaborate what our findings and the existing evidence imply for theory and policy. We start by discussing the existing evidence on the associations between reward distributions, risk preferences and errors. Then, we shift our attention to the evidence on how cognitive capacity affects risk attitudes and decision errors.

4.1 Distribution of Prizes

The idea that utility should optimally adapt to the distribution of prizes in a chooser's environment is not new (Robson 2001, Rayo and Becker 2007, Glimcher 2010, Woodford 2012). A small number of theoretical papers have demonstrated conditions under which the S-shaped utility function, as proposed in Prospect Theory is optimal if the chooser's goal is to minimize decision errors or maximize expected earnings (see (Vieider 2025) for review). While prior research has mainly aimed to justify the empirically observed S-shaped utility function, our work is the first to explicitly investigate how the distribution of prizes in the environment shapes the curvature of an optimal utility function. This enables us to bridge our normative model to observed behavioral patterns, providing a deeper understanding of the systematic variation across people.

4.1.1 Distribution of Prizes and Risk Attitudes

We predict a sharp difference in utility curvature depending on whether choosers are exposed to positively or negatively skewed prize distributions. Specifically, we show that optimal choosers who more frequently encounter lower-valued prizes will develop concave utility functions, while those who are exposed to higher-valued prizes will exhibit convex utility, even if the objective value range of the prizes remains the same. These theoretical predictions align with a robust body of empirical evidence linking risk attitudes and wealth. Lower-wealth households exhibit behaviors that are consistent with more concave utility functions. They are significantly less likely to participate in the stock market, to their financial detriment (Calvet *et al.* 2007). This wealth-risk correlation has been documented in large-scale representative surveys (Falk *et al.* 2018) and in laboratory experiments (Tymula *et al.* 2013). Similarly, lower willingness to take risk has been usually observed among women which could be driven by their more frequent exposure to negatively skewed financial payoffs and thus lower financial payoff expectations (Levy *et al.* 2025). The fact that gender differences in risk attitudes are not universally observed lends further support to the hypothesis that these differences are adaptive responses to the economic environment.

In parallel, laboratory studies have shown that indeed utility functions are malleable and responsive to the distribution of past rewards both for risky and riskless choice. For example, Khaw et al. (2017) and Guo & Tymula (2021) found that even just brief exposure to low versus high prize distributions (under riskless conditions and without financial consequences) leads to reliable shifts in utility, as predicted in our model. Field studies echo these findings: people who entered the job market during a recession tend to invest less in stocks over their lifetimes (Malmendier and Nagel 2011), and those with lower income and socioeconomic status expect worse financial returns from the same financial opportunities and as a result take fewer financial risks and make less money (Das *et al.* 2020).

4.1.2 Distribution of Prizes and Decision Errors

The slope of the utility function determines the discriminability between the options and hence the likelihood of decision errors. Our model predicts that for choosers aiming to maximize earnings, it is optimal to exhibit steeper utility (and hence better discriminability) for prizes that are encountered more frequently and contribute more substantially to higher earnings. This insight allows our framework to account for a range of the well-documented behavioral phenomena, including many aspects of the curse of choice (i.e. the increase in decision errors as the number of prizes increases in a way that violates the regularity axiom), violations of the regularity axiom outside the curse of choice, and systematic changes in utility slope as prize distributions widen.

When applied to socioeconomic differences, our model our model offers a compelling explanation for why people with lower incomes often perform well in low-stakes decisions but struggle in higher-stakes scenarios(Bertrand *et al.* 2004, Shah *et al.* 2012, Mani *et al.* 2013). Consistent with this literature, we conclude that people living in conditions of poverty or abundance are tuned to the statistical structure of their environments which is then reflected in their decisions.

4.1.3 New Predictions and Policy Implications

Our framework also generates several novel predictions that, to the best of our knowledge, have not yet been empirically tested. We find that increasing either the number of prizes in the environment or the number of options in the choice set, each independently can make the S-shaped utility functions be more likely to be optimal—even when the underlying prize distribution is not normal. This untested prediction could explain why so many experiments observed S-shaped utility functions despite using skewed rather than normally distributed rewards.

We find that adding more options to the choice set increases the steepness of the utility for the higher-value prizes. Intuitively this follows from the fact that as the choice set size grows, it becomes more likely to include high-value prizes drawn from the environmental prize space, making it increasingly important to distinguish between them. In contrast, lower-value prizes become less relevant as they are increasingly unlikely to be selected as choice set sizes grow. This prediction offers another promising explanation for context-dependent changes in utility curvature to be tested in future research.

Overall, the evidence that the slope and curvature of utility function adjusts to chooser's prize distribution is strong. From a policy perspective this implies that people who are adapted to different prize distributions will make different decisions when faced with the same opportunities. These differences in choices may reflect their environmental conditions rather than their fixed preferences. It is thus important that policy makers are aware of these mechanistic dependencies that are likely built into our nervous system. Providing equal opportunities without recognizing the impact of heterogenous environments on decisions may be insufficient. Our approach illustrates that even if provided with equal opportunities, choosers with utilities optimized for different environments would make different choices. This suggests that optimizing both the choice environment and the methods by which we communicate the choice environment to choosers may be beneficial. Studies like (Frydman and Jin 2021, Guo and Tymula 2021, Khaw *et al.* 2021) highlight that it may be possible to adjust people to decision-relevant prize distributions within quite short timeframes before they are asked to make a choice. This suggests that, cleverly designed, even inconsequential, short exposures to some decision scenarios and other such decision-making aids could help people achieve outcomes more aligned with their true objectives.

4.2 Cognitive Constraints

In addition to prize distribution, the other primitive that influences utility curvature in our framework is cognitive capacity. We predict that as the cognitive capacity becomes less constrained, linear utility functions emerge. However, as the capacity constraint becomes more binding, utility curvature becomes environment-dependent: in positively skewed environments, optimal utility becomes concave; in negatively skewed environments, it becomes convex. In our model, cognitive capacity is determined by neurobiological constraints on computational precision, the intrinsic noise associated with the physical instantiation of utility representations in the brain. For decades, neurobiologists have examined the relationship between brain structure and representational precision and have concluded, unsurprisingly, that smaller brains, fewer neurons, and lower metabolic capacities are all associated with lower precision. Gray matter volume (our proxy for capacity) is relatively stable in the short term but reduces over the lifespan of an adult. Thus, we predict that one important source of variation in behavior should happen along the life course developmental trajectory of the human brain. We know, for example, that the number of neurons in the human brain begins to decline around the age of 25, initially gradually and then more steeply after the mid 40s. These declines in neuron numbers are also correlated quite tightly with decreasing cognitive capacity. Below we review the existing evidence on how risk attitudes and decision errors change across lifespan and their direct association with the gray matter volume.

4.2.1 Cognitive Constraints and Risk Attitude

Our finding that utility curvature is associated with cognitive capacity suggests that when the environment remains unchanged, it should be a stable individual trait. Supporting this claim, previous studies (Gilaie-Dotan et al., 2014) found that risk attitudes correlate with gray matter volume in the posterior parietal cortex. People who have more neurons in this area (and thus can produce more action potentials, an equivalent of higher cognitive capacity in our model) are more risk neutral. Grubb et al. (2016) have shown that the age-related decrease in gray mattery volume in the posterior parietal cortex is associated with more concave utility functions. Moreover, it is the age-related decrease in gray matter volume and not the chronological age per se that drive the increase in concavity in ageing. Our linkage of neural/cognitive precision with cognitive capacity sheds light on the origins of these observed life course effects on the shape of the utility function.

Our findings also align well with the observation that older adults exhibit greater risk aversion (more curvature) than do younger choosers, at least in the domain of gains (e.g., Barsky et al., 1997). Additionally, studies examining risk taking across both gains and losses (Tymula *et al.* 2013) have shown that while older adults are more risk averse for gains, they are more risk seeking for losses than younger adults. This pattern aligns with our prediction that lower cognitive capacity leads to greater deviations from risk neutrality.

Outside the aging literature, a growing body of evidence indicates that lower socioeconomic status is associated with reduced gray matter volume across the life span (Noble *et al.* 2015) and a steeper rate of age-related decline (Steffener *et al.* 2016). This may help explain the lower financial risk-taking observed in lower socioeconomic status groups, and faster increase in risk aversion with age within these populations (Schurer 2015).

4.2.1 Cognitive Constraints and Decision Errors

There is plenty of literature that associates aging with a decline in cognitive capacity including more costly decision errors. Chung *et al.* (Chung *et al.* 2017) measured individual study participants' gray matter volume and the propensity to violate the Generalized Axiom of Revealed Preference violations in older adults. They found that the age-related violations can be traced to the decline in gray matter volume. Behavioral studies, for example (Tymula *et al.* 2013), showed that older adults (over 65 years old) are much more likely to choose stochastically dominated lotteries than younger adults. These mistakes were not a result of the lack of understanding of the task but rather they were fully consistent with the implications of our definition of capacity. When choosing between $5 for sure and a lottery that pays $5 with some probability, older adults were more likely to choose the first-order stochastically dominated lottery the higher was the probability with which it would pay $5, that is the closer the two options were in their expected value. The decisions of older adults in this study, were also generally more random.

4.2.1 Policy

From a policy perspective, the good news is that gray matter volume is not fixed and, like muscle function, can improve with use. For example, studies have shown that London taxi drivers develop increased hippocampal volume due to intensive spatial navigation training (Maguire *et al.* 2006). These findings suggest, it is not far-fetched to imagine cognitive or behavioral interventions that promote structural brain changes, thereby enhancing or preserving financial decision-making abilities across the lifespan.

5. Conclusions

Our paper advances a novel framework for understanding human decision-making by emphasizing the importance of the probabilistic structure of the reward environment, acting as a *prior distribution* from which decision-makers must infer how to achieve maximization given their inherently limited precision in encoding the value of prizes. From the perspective of our framework, any choice problem that does not explicitly specify the distribution of potential prizes and the limited capacity to encode value remains underspecified from the standpoint of a fully usable theory of choice. This marks a conceptual departure from standard approaches, which often describe prizes without specifying their associated probabilities and do not use cognitive capacity as a primitive.

Drawing inspiration from Pascal, like some other recent papers (see Vieider (2025) for review) we adopt the premise that decision-makers aim to maximize the objective value of the rewards that they receive. In contrast to many traditional models, we abandon preferences as the primitives of our model. Curved utility functions emerge as a consequence of cognitively constrained decision-makers maximizing rewards in a deterministic choice setting. Thus, our findings can well-explain puzzling findings like those in Oprea (Oprea 2024) who found curved utility functions measured in riskless choice tasks correspond to utility curvature measured in the same decision makers exhibit in risky choice. Our framework is also in line with the findings in Barretto-García et al. (2023) who found a close link between number representations in riskless settings and willingness to take risk.

Consistent with prior theories, we assume that individuals transform objective values into subjective ones via a measurement utility-like function and then compare them to chose the one that yields the highest utility. However, our model diverges by endogenizing this transformation: the utility function is treated not as a fixed but rather as a choice variable under the control of the decision-maker who wants to maximize their long-run earnings in a given prize environment. This allows us to separate the subjective representational structure of value from the objective function itself, offering a richer account of individual differences.

Another central feature of our framework is the role of cognitive capacity, the precision with which the subjective representation of the objective value of rewards is measured. We initially treat precision as an exogenous but variable parameter and explore the space of optimal utility functions across different precision levels. We find that individual variation in choice behavior is driven by differences in cognitive capacity, in addition to differences in prize distributions. We extend our framework by considering the possibility that cognitive capacity itself may be subject to optimization. This formulation allows us to reframe the decision-making problem as an optimization problem where the only free parameter is the cost of cognitive capacity. Perhaps surprisingly, we observe that the marginal benefits of increased cognitive capacity are modest suggesting that the returns to capacity may often be outweighed by even small associated costs.

We note that in the companion psychology paper Glimcher, Singha and Tymula (2025), we examine an alternative objective function: the minimization of decision errors rather than the maximization of returns. We find that this alternative perspective yields behaviorally similar predictions in many—but not all—environments, suggesting that while different objectives can produce overlapping behavioral patterns, they may diverge under certain conditions. This comparative analysis underscores the importance of clarifying the assumed goals of the decision-maker when constructing predictive models of choice.

One potential limitation of our approach is that we derive the optimal riskless and atemporal utility functions. Several papers provided evidence that utility elicited under riskless and risky conditions differ (Andreoni and Sprenger 2012, Chung *et al.* 2019) but others found that they are surprisingly alike (Oprea 2024). Given that this debate is still ongoing, although the implications of our riskless utility functions for risky choice are convincing (see section 4 for empirical evidence), we caution the reader to treat these extensions with caution. We note that we are in the process of generalizing our approach risky choice. This requires either assuming the independence axiom (effectively assuming no limits to the precision for representing probability) or extending our approach to the domain of probabilities. Our approach could also in principle be generalized to intertemporal choices by assuming an intertemporal objective function choosers seek to maximize. Under that set of conditions, one could analytically and/or numerically derive optimal discount functions in much the same way that we derive optimal utility functions here. All of these are important approaches that require further examination in future work.

Bibliography

Andreoni, J. and Sprenger, C., 2012. Risk Preferences Are Not Time Preferences. *American Economic Review*, 102 (7), 3357–3376.

Barretto-García, M., de Hollander, G., Grueschow, M., Polanía, R., Woodford, M., and Ruff, C.C., 2023. Individual risk attitudes arise from noise in neurocognitive magnitude representations. *Nature Human Behaviour*, 7 (9).

Barsky, R.B., Juster, F.T., Kimball, M.S., and Shapiro, M.D., 1997. Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study. *Quarterly Journal of Economics*, 112 (2).

Bernoulli, D., 1954. Exposition of a New Theory on The Measurement of Risk. *Econometrica*, 22 (1), 23–36.

Bertrand, M., Mullainathan, S., and Shafir, E., 2004. A behavioral-economics view of poverty. *In*: *American Economic Review*.

Bucher, S. and Brandenburger, A., 2022. Divisive Normalization is an Efficient Code for Multivariate Pareto-Distributed Environments. *Proceedings of the National Academy of Sciences*, 119 (40).

Calvet, L.E., Campbell, J.Y., and Sodini, P., 2007. Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy*, 115 (5).

Chung, H.-K., Glimcher, P.W., and Tymula, A., 2019. An Experimental Comparison of Risky and Riskless Choice – Limitations of Prospect Theory and Expected Utility Theory. *American Economic Journal: Microeconomics*, 11 (3), 34–67.

Chung, H.-K., Tymula, A., and Glimcher, P., 2017. The Reduction of Ventrolateral Prefrontal Cortex Gray Matter Volume Correlates with Loss of Economic Rationality in Aging. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 37 (49), 12068–12077.

Das, S., Kuhnen, C.M., and Nagel, S., 2020. Socioeconomic Status and Macroeconomic Expectations. *Review of Financial Studies*, 33 (1).

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., and Sunde, U., 2018. Global evidence on economic preferences. *Quarterly Journal of Economics*.

Friedman, D., 1989. The S-Shaped value function as a constrained optimum. *The American Economic Review*, 79 (5).

Frydman, C. and Jin, L.J., 2021. Efficient coding and risky choice. *Quarterly Journal of Economics*.

Gilaie-Dotan, S., Tymula, A., Cooper, N., Kable, J., Glimcher, P., and Levy, I., 2014. Neuroanatomy predicts individual risk attitudes. *Journal of Neuroscience*, 34 (37).

Glimcher, P., 2010. *Foundations of Neuroeconomic Analysis*. Oxford University Press.

Glimcher, P., 2022. Efficiently irrational: illuminating the riddle of human choice. *Trends in Cognitive Sciences*, 26 (8), 669–687.

Grubb, M., Tymula, A., Gilaie-Dotan, S., Glimcher, P., and Levy, I., 2016. Neuroanatomy accounts for age-related changes in risk preferences. *Nature Communications*.

Guo, J. and Tymula, A., 2021. Waterfall illusion in risky choice – exposure to outcome-irrelevant gambles affects subsequent valuation of risky gambles. *European Economic Review*, 139, 103889.

Heng, J.A., Woodford, M., and Polania, R., 2020. Efficient sampling and noisy decisions. *eLife*, 9.

Kahneman, D. and Tversky, A., 1979. Prospect Theory. *Econometrica*, 47 (2), 311.

Khaw, M.W., Glimcher, P.W., and Louie, K., 2017. Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences*, 114 (48), 12696–12701.

Khaw, M.W., Li, Z., and Woodford, M., 2021. Cognitive Imprecision and Small-Stakes Risk Aversion. *Review of Economic Studies*, 88 (4).

Lennie, P., 2003. The Cost of Cortical Computation. *Current Biology*, 13 (6), 493–497.

Levy, J., Ru, O., and Tymula, A., 2025. *The Gender Reference Point Gap: Evidence from a Representative Sample*.

Maguire, E.A., Woollett, K., and Spiers, H.J., 2006. London taxi drivers and bus drivers: A structural MRI and neuropsychological analysis. *Hippocampus*, 16 (12), 1091–1101.

Malmendier, U. and Nagel, S., 2011. Depression Babies: Do Macroeconomic Experiences Affect Risk Taking? *Quarterly Journal of Economics*, 126 (1), 373–416.

Mani, A., Mullainathan, S., Shafir, E., and Zhao, J., 2013. Poverty impedes cognitive function. *Science*, 341 (6149).

McFadden, D., 1974. Conditional logit analysis of qualitative choice behavior. *In*: P. Zarembka, ed. *Frontiers of Econometrics*. New York: Academic Press.

Netzer, N., 2009. Evolution of Time Preferences and Attitudes toward Risk. *American Economic Review*, 99 (3), 937–955.

Netzer, N., Robson, A., Steiner, J., and Kocourek, P., 2025. Risk perception: measurement and aggregation. *Journal of the European Economic Association*.

Noble, K.G., Houston, S.M., Brito, N.H., Bartsch, H., Kan, E., Kuperman, J.M., Akshoomoff, N., Amaral, D.G., Bloss, C.S., Libiger, O., Schork, N.J., Murray, S.S., Casey, B.J., Chang, L., Ernst, T.M., Frazier, J.A., Gruen, J.R., Kennedy, D.N., Van Zijl, P., Mostofsky, S., Kaufmann, W.E., Kenet, T., Dale, A.M., Jernigan, T.L., and Sowell, E.R., 2015. Family income, parental education and brain structure in children and adolescents. *Nature Neuroscience*, 18 (5).

Oprea, R., 2024. Decisions Under Risk are Decisions Under Complexity. *American Economic Review*, 12 (114), 3789–3811.

Polanía, R., Woodford, M., and Ruff, C.C., 2019. Efficient coding of subjective value. *Nature Neuroscience*, 22 (1).

Rayo, L. and Becker, G.S., 2007. Evolutionary Efficiency and Happiness. *Journal of Political Economy*, 115 (2), 302–337.

Robson, A.J., 2001. The Biological Basis of Economic Behavior. *Journal of Economic Literature*, 39 (1), 11–33.

Robson, A.J., Whitehead, L.A., and Robalino, N., 2023. Adaptive utility. *Journal of Economic Behavior and Organization*, 211.

Schurer, S., 2015. Lifecycle patterns in the socioeconomic gradient of risk preferences. *Journal of Economic Behavior and Organization*.

Shah, A.K., Mullainathan, S., and Shafir, E., 2012. Some consequences of having too little. *Science*, 338 (6107).

Steffener, J., Habeck, C., O'Shea, D., Razlighi, Q., Bherer, L., and Stern, Y., 2016. Differences between chronological and brain age are related to education and self-reported physical activity. *Neurobiology of Aging*, 40.

Steverson, K., Brandenburger, A., and Glimcher, P., 2019. Choice-theoretic foundations of the divisive normalization model. *Journal of Economic Behavior & Organization*, 164, 148–165.

Stewart, N., Reimers, S., and Harris, A.J.L., 2015. On the origin of utility, weighting, and discounting functions: How they get their shapes and how to change their shapes. *Management Science*.

Tymula, A., Belmaker, L.A.R., Ruderman, L., Glimcher, P.W., and Levy, I., 2013. Like cognitive function, decision making across the life span shows profound age-related changes. *Proceedings of the National Academy of Sciences of the United States of America*, 110 (42).

Vieider, F., 2025. *Neuro-Biological Origins of PT functionals*.
Woodford, M., 2012. Prospect Theory as efficient perceptual distortion. *American Economic Review*, 102 (3), 41–46.