

Variability in Complex Constructs: Inferring Risk Preference and Temporal Discounting

Nidhi V. Banavar* (nbanavar@uci.edu) & Aaron M. Bornstein (aaron.bornstein@uci.edu)

We examine the prevalence and extent of variability across measurements of supposedly stable behavioral economic traits. We begin by reviewing how these traits are conceptualized in behavioral economics, and how different instruments for eliciting them lead to variability in their measurements. We then consider factors such as experiment structure, affect, and context, known to influence or correlate in some way with the inferred values of these constructs: from domain or “subject-level” influences to local influences. We introduce the idea that an important – and cognitively meaningful – source of potential variation in experimentally-inferred measures may come from temporal sequence or the influence of trial order. Finally, we discuss how some of these sources of variation may not be ultimately all be brought under experimental or analytical control, and propose that they should instead be exposed and considered for their predictive value in different settings.

Introduction

In this paper, we provide a synoptic perspective on two widely studied psychological constructs: risk preference (how individuals behave under known and immediate uncertainty) and temporal discounting (how individuals behave under future uncertainty). Increasingly, important decisions about an individual are being informed by an assessment of how that individual makes decisions under uncertainty. These can be high-stakes interventions (e.g. personalized vaccine incentives - Andreoni et al, 2016), but the measures they depend upon are known to be variable across contexts (Peters & Büchel, 2011) and may have limited predictive power for real-world consequences. We discuss how the different methods used to measure these constructs are subject to distinct forms of variability, each of which themselves can be valuable in different settings. Some of these measures may be more robust to time, and others may be more robust to framing

*to whom the correspondence may be addressed. Funding for this project was provided by NIA R21AG072673 (to AMB) and by the Brain and Behavior Research Foundation NARSAD Young Investigator Award (AMB).

effects. Those of us who study and apply risk and temporal discounting measures would therefore benefit not only from the use of appropriate methods but also from an understanding of their endemic variability.

Behavioral Economic Definitions

To understand the nature of variability in the constructions of interest, we must first define the theoretical framework under which they are measured, neoclassical behavioral economics.

Rationality. A fundamental assumption in many behavioral economic models of choice behavior under uncertainty is that ‘rational’ individuals act in order to maximize their expected utility (von Neumann & Morgenstern, 1944). That is, assuming that they meet the axioms demarcating rationality, a utility function can be well-defined (von Neumann & Morgenstern, 1944). Individuals can then be characterized as acting in order to maximize these utility – or satisfaction – functions. This is known as Expected Utility Theory (EUT). Other classical theories of utility such as Random Utility Theory which assumes stochasticity (Thurstone, 1927) and Subjective Expected Utility Theory (SEUT) which extends EUT as implied by its name (Savage, 1954). Critically, economic rationality necessitates consistency: some inherent stability or structure in preference and choice behavior. Much empirical evidence and theoretical skepticism, however, has cast doubt on the value of the rational (ideal) individual, revealed preference, and how closely the behavior of said individual matches actual human behavior –famously violated in the Allais Paradox (Allais, 1990), Ellsberg Paradox (Ellsberg, 1961) and in framing effects (Tversky & Kahneman 1979). Despite all this, it is for good reason that (S)EUT has continued to dominate the behavioral economic study of decision making under uncertainty: it is a compelling normative framework that has been used across a wide range of domains including choice behavior, public policy, and medicine (Weber, 2010). We therefore consider more carefully the formalization of decision-making under uncertainty within this framework.

Risk. Immediate uncertainty can be in the form of risk – where probability distributions are fully known – or ambiguity – where the probabilities are fully or partially unknown. For example, an individual making a choice between a 25% chance of \$20 (and 75% chance of \$0) or \$5 guaranteed, is making a choice under risk. Alternatively, if the subject is choosing between the guaranteed \$5 and a 25% chance of \$20 without knowing what the remaining 75% would give them, they are deciding under partial ambiguity – they have

incomplete information about the underlying probability distribution. While conceptually intimately related, people behave differently under risky compared to ambiguous circumstances, with evidence for even more nuanced differences in clinical populations (e.g. Konova et al, 2020). As quotidian decision making rarely involves complete probabilistic information, decision making under ambiguity is of considerable ecological interest. Due to the limited scope of this paper, however, we will primarily consider decision making under risk.

Expected Utility, then, is usually characterized as a power function: $Expected\ Utility = p \cdot v^\alpha$, where p represents the probability of a given reward and v represents the objective reward (e.g. dollar amount of reward). The curvature (α) of the expected utility function, known as risk tolerance, has very specific interpretations: if an individual is risk averse, they prefer less uncertain outcomes even if the dollar reward is lower. This is characterized by a concave utility function, with $\alpha < 1$ (individuals whose utility function is convex with $\alpha > 1$ are classified as risk seeking and $\alpha = 1$ as risk neutral). Famously, Kahneman and Tversky showed that these characterizations depend on a reference point (the domain in which decisions are made) – they held if people were making choices between rewards they could possibly win, but flipped if people were making choices between rewards they could possibly lose (i.e. risk aversion in the loss domain yields a convex utility function) (Cumulative Prospect Theory (CPT); Kahneman & Tversky, 1979). As there are many other ways to parametrize an individual’s decision making under risk (see *Delineating Risk*), we use risk tolerance (or α) to reference the parameter predicated on EUT/CPT (and variants) and risk preference to refer to the more general latent variable/concept. We further note the related phrase risk perception, which is often confounded with risk tolerance and preference, is generally defined as the recognition of inherent risk – the ability, given internal and external circumstances, to appropriately assess the riskiness of a situation (Hunter, 2002). This is contrasted with risk tolerance or preference, as they deal more specifically with an individual’s willingness to engage in risky decision-making, usually when available options are of equal expected value.

Temporal Discounting. Models of temporal discounting study the interaction between time, value, and uncertainty and aim to capture how individuals attribute differential weight to choice options closer to the

present compared to the (distant) future. A rational, and therefore consistent, individual would use an exponential discounting function (Samuelson 1937): $Utility = v \cdot e^{-kd}$, where v again represents the objective amount of reward; d the delay, or how far away from the present the reward would be received and k the discount factor. This is also known as a linear exponential model as a linear utility function (with $\alpha = 1$, interpreted as discussed in the previous paragraph) is assumed. Exponential discounting indicates that the rate at which an individual discounts is constant over time. However, decades of empirical work has demonstrated that this does not explain human behavior well, and that the linear hyperbolic discounting model is a significantly better fit: $Utility = \frac{v}{1+kd}$, where all variables maintain their interpretations (Thaler 1981, Kable & Glimcher, 2007). This model also assumes that individuals are risk neutral, but critically allows time-varying discount rates: steeper discounting when the delayed option is closer in time. The discount factor can range from 0 to 1 – where a value of 0 would mean that an individual considers only the dollar amounts offered regardless of temporal distance. Interestingly, animals tend to be even more myopic in their discounting (Loewenstein et al, 2015).

Just as with risk, there are many ways to parametrize how an individual's discounting decays (see *Delineating Temporal Discounting*). We use the term discount factor (or k) to reference parameters predicated on EUT (and variants), and the more general temporal discounting to reference the concept. Finally, we note that an individual's predilection towards selecting the option closer in time at the cost of an objectively better (as per economic rationality) reward is also commonly referred to as their impulsivity – individuals with high ks are usually interpreted to be more impulsive. Like risk tolerance, impulsivity has also been used to mark differences in clinical populations (e.g. Addiction: Bickel et al, 2014; Major Depressive Disorder: Pulcu et al, 2015). Finally, we note that impulsivity is a multidimensional construct where the dimensions themselves are yet to be agreed upon by researchers (Evenden, 1999), and therefore should not strictly be used interchangeably with temporal discounting.

Relating the two. While risk tolerance and discount factor are distinct in both their psychological and economic formalizations, they are intimately related. Indeed, they must be by definition at least within the

Expected Utility frameworks, as both parameters explicitly invoke Subjective Value. The standard (exponential/hyperbolic) discounting model implicitly assumes that a given individual is risk neutral, and this imposes a systematic bias on measures of discount factor (Lopez-Guzman et al, 2018). That such a bias exists corresponds with a well-established empirical finding that people are generally risk averse in the gain domain and risk seeking in loss (Tversky & Kahneman, 1989); they typically aren't risk neutral. They are also conceptually related as they both involve notions of uncertainty (immediate vs. temporal/delayed). Other theories that explicitly link risk preference and temporal discounting include Construal-Level Theory (Lieberman & Trope, 2003) which posits that an individual's change in (subjective) value is due to changing mental representations and a differential focus on concrete vs abstract features, informing the manner in which an individual both discounts future reward and behaves under risk (Leiser et al, 2008). Dual-process models of deliberation and affect suggest that risky decision making and intertemporal choice (ITC) may also be linked in how people trade off the desirability of presented options with the cost of willpower or effort required (Loewenstein et al, 2015).

Empirical studies in animals have shown that animals are sensitive to the frequency with which they must make risky choices: one study found that rhesus monkeys that choose between risky gambles every 3s demonstrate risk seeking behavior, while others in rats and birds found evidence for risk aversion when choices were made approximately every 30s (Hayden and Platt, 2007). More concretely, Hayden and Platt found that rhesus monkeys preferred certain options over risky ones with increasing delay (larger inter-trial intervals) between choices (Hayden and Platt, 2007). That an ostensibly simple and unrelated manipulation in how animals, assuming some degree of comparative equivalency, choose between gambles in an experiment leads to behavior with distinct interpretations suggests a relationship between time, uncertainty, and choice that is evolutionarily old. This is consistent with a 2013 meta-analysis on neuroscience data sets searching for a neural value system – the two primary regions identified being the (evolutionarily newer) ventro-medial prefrontal cortex and the (older) ventral striatum (Bartra et al, 2013). There are also other evolutionarily older ways in which neural information processing and representation could relate risky decision making, intertemporal choice, and cognition more broadly (see *Trial-Trial*

Temporal Dependence). We note that despite the fact that these constructs are related across many dimensions, they are usually studied separately (i.e. risky decision-making experiments vs. intertemporal choice experiments).

Rationality Revisited. Behavioral economic parameters of interest, useful in both explanatory and predictive capacities, are multi-dimensional constructs that also manifest differently at different time scales (see *Delineating Risk, Delineating Temporal Discounting*). There have accordingly been decades of debate over whether it is appropriate to conceptualize these constructs as trait-like variables, with strong evidence of both stability and variability in choice behavior and subsequent inferred parameter values. Choice behavior, through which these constructs are regularly quantified, however, is subject to a wide range of influences (see *Delineating Risk, Delineating Temporal Discounting, Variability in Experiments*). We might consider an individual in a financial difficult circumstance, or more mundanely, trying to purchase an out of budget treat for themselves: they may be biased in an intertemporal choice experiment towards more immediate rewards, and therefore be characterized as very impulsive. In that their current desire for the immediate reward is identified, measurements and standard interpretations of temporal discounting are appropriate. However, it is an entirely different, but related question as to whether this measurement is reflective of their ‘innate’ impulsivity – if such a construct is even meaningful (reasonable test-retest reliability, at the very least, suggests that could be, e.g. Kirby, 2009; Frey et al, 2017). However, it is difficult to make a decisive statement on this, after all, individuals are asked explicitly to act in accordance with their true preference. In fact, if individuals were truly compliant with task instructions, there would be little dissociation between (personal) circumstantial and task-congruent influences. Similar logic holds true for individuals participating in a risky decision-making task – there are a multitude of reasons why an individual behaves the way they do in any given moment. The relationship between goals, especially clearly articulated goals, context, behavior, and parameters of interest is one that requires careful thought. This is especially the case when more “extreme” behaviors are usually interpreted as deviations from rationality and its subsequent throng of implications.

The wide range of observed behaviors that deviate from economic rationality, then, are usually either characterized as “not as irrational as they may seem,” or as necessitating a new conceptual, veridical framework (Vlaev, 2018). Recently, Vlaev formalized a theoretical compromise, leveraging the definition of rationality most prominent in cognitive science (Anderson, 1990): it is more appropriate to consider human behavior as locally rational – that humans make (environmentally) contextualized rational inferences as opposed to universal (Vlaev, 2018). Given various constraints in human decision making (e.g. limits on information processing) and the sheer depth of computation theoretically necessary for globally consistent choice (Simon, 1990), this strikes as a more plausible casting of the human implementation of [economic] rationality. This framework therefore implies that value itself is not well defined or consistent. That is, as others (e.g. Slovic 1995) have proposed, preference is constructed and not just revealed during elicitation (though note that this does not mean that preference is necessarily constructed from scratch, or independently of previous experiences, each time). Thus, Vlaev synthesizes that, using (limited) resources and privileged information, value, local comparisons, and subsequent inferences are all computed online and in a rational, sequential manner. Despite suggesting a lack of consistency in value and its predicated constructs like risk tolerance and discount factor, we do not argue that the notion of stability in these constructs should be entirely divorced from their conception. It is, after all, impossible to formulate constructs that are sufficient in depth and breadth.

The remainder of the paper is outlined as follows. First, we discuss variability in eliciting and measuring both risk preference and temporal discounting. Then we consider different ways in which we can analyze experimental data and relate these measures to preference. Finally, we consider the ideas of task-incongruent temporal dependencies (i.e. perseveration and serial dependence).

Delineating Risk

We first consider the different ways in which risk preference can be measured behaviorally.

Measurements. Experiments designed to estimate an individual’s risk preference generally fall into one of three categories: statistically dependent sequential choices (SDSC, e.g. Balloon Analogue Risk Task, n-

armed bandit tasks); statistically independent ordered choices (SIOC, e.g. Holt and Laury gambles) and statistically independent single choices (SISC, e.g. lottery tasks) (Pedroni et al, 2017). Other features that can vary across tasks include how choices are displayed (numerically – with monetary values and probabilities listed, graphically – with graphical or pictorial depictions of probability, or both); choice domain (gain, loss, or mixed); incentivization (e.g. Becker-deGroot-Marschak random draw, cumulative reward); the presence or absence of feedback (that is, the immediate realization of their choice resulting in feedback informing them of their win/loss) and the amount of time an individual has to respond (Pedroni et al, 2017). While SDSC tasks often explicitly model learning and other possible temporally evolving processes and dependencies, SISC tasks focus on “in the moment” decision making and typically consist of randomized, and therefore temporally unstructured choice sets. A further, related, distinction across these experiments is the “description – experience” gap: that people behave differently when complete information is provided about the problem (and by extension the environment) versus when they are provided incomplete information, and need to rely on experience (previous or current) (Hertwig & Erev, 2009). If we consider description – experience, risk / ambiguity and the presence or absence of feedback, we can further taxonomize these experiments. SDSC tasks are typically experiential, while SISC tend to be descriptive, with feedback acting as an important arbitrator between SDSC and SISC and between whether the individual is making decisions under ambiguity or risk. In SDSC tasks, incomplete information about the probabilistic structure of the environment (or bandit machines, for example) is reducible – people can actively learn about and mitigate the underlying uncertainty through feedback (usually in the case of rewards won or lost after a choice). In SISC tasks, however, the reducibility of ambiguity is entirely dependent on the construction of the choice set and the presence of feedback. In fixed ‘unstructured’ (i.e. choice set does not change over the task like with staircasing) experiments without feedback, the underlying uncertainty is irreducible. The individual makes choices in the dark and with, in theory, only their preference and description of the problem to guide them. These are descriptive decisions under both risk and ambiguity, as information is explicitly presented to the individuals, with nothing to be “learnt” as lotteries presented are usually fixed (i.e. at least 25%, 50%, 75% chances of winning). Similarly structured

experiments that involve feedback, however, can allow for individuals to learn about what the underlying probability of the various gamble types presented, much like in bandit tasks, except that subsequent choices are unrelated to each other. Feedback therefore allows for individuals to “experience” the consequences of their decisions and, depending on the goal of the experiment, there is variability in how feedback is expected to influence trial and aggregate choice behavior (Barron & Erev, 2003, Brooks & Sokol-Hessner, 2020). Usually, however, the standard modeling framework of SISC experiments, especially in the context of inferring these parameters of interest, is to treat data as explicitly descriptive and not account for potential transient within-task influences or learning, however task-irrelevant they might be.

Many risk preference elicitation methods exist in the literature. Beyond behavioral experiments, there exist more subjective measures, usually in the form of Likert scales (e.g. “how risk seeking are you in general?”) or surveys, such as the Domain Specific Risk Taking (DOSPERT) Scale (Weber et al, 2002, see Frey et al, 2017 for a comprehensive list). Finally, measures of an individual’s risk preference can also come from frequency measures by tabulating the occasions on which an individual engages in risky behavior, though we note that this information is usually also collected through self-report (Frey et al, 2017) unless, for example, in a clinical setting.

Variability. That risk preference is complex is intuitive, if not patent in the many means by which it can be defined and assessed. Changes in risk preference have been observed as a function of affect/motivational state (fear increases risk aversion while anger decreases it: Kugler et al, 2012); age (older adults are more risk averse: Tymula et al, 2012); clinical disorders (patients with substance use disorder are more risk tolerant: Konova et al, 2020); and sex (women are more risk averse: Croson & Gneezy, 2009). We note that for most studies that find structured evidence of the malleability of risk preference, there are studies that find no evidence of any such sensitivity (e.g. no systematic effect of stress: Sokol-Hessner et al 2016). Risk preference is also thought to be related to important variables such as income, intelligence, and education – though a recent study in a large (N = 916) diverse cohort of adults found that only sex and age have robust, consistent associations with risk preference (Frey et al, 2020). Importantly, this study found that the relationship between risk preference and these correlates varied as a function of how risk preference was

measured, with subjective measures being more sensitive to these correlates relative to experimental measures. Earlier work by the same group sought to examine whether these ostensibly different measures of risk could be consolidated into a single latent variable R , much like g (intelligence) (Frey et al 2017). Using 37 different risk elicitation measures in a sample of 1507 from two different countries, the authors found that they were indeed able to extract a temporally stable R that accounted for 50% of observed variation in a factor analysis. Critically, however, almost all this stability was attributed to measures elicited from surveys (e.g. DOSPERT) and frequency counts of risky behavior, also measured through surveys. Further, subjective and frequency measures had much higher temporal stability and correlations between and within themselves relative to measures elicited experimentally (ranging from SDSC to SISC tasks). Nonetheless, recent work has demonstrated the value of temporal fluctuations, finding in a clinical setting that only week-to-week fluctuations in experimentally elicited measures of ambiguity tolerance and recent risky behavior (e.g. recent drug use) were predictive of future real-world behavior under uncertainty (Konova et al, 2020). More generally, other studies have also found relatively low correlations between experimentally induced measures of risk tolerance, including differential contextual or emotional sensitivities (Guan et al 2020, Kugler et al 2012, Milroth et al 2020, Pedroni et al, 2017, Radulescu et al 2020, Sokol-Hessner et al 2016). As risk tolerance is well established, and by definition subjective and relative (Weber, 2004), and experiments themselves can vary widely in their construction, it is perhaps unsurprising to find such high levels variability in behavioral experiments – as self-reports might manifest more like personality traits than as functions of socioeconomic status or cognitive ability (Frey et al 2017). Self-reports also assess risk preference at a different (global) timescale, and therefore elicit qualitatively different information. Further, as the questionnaires often ask individuals to respond hypothetically and “in general,” it would be more appropriate to characterize these measures as decision making under ambiguity, not risk. It might be even more appropriate to consider these measures as meta cognitive: that they reflect an individual’s thinking about how they think about the question vs their thinking during actual choice. Local rationality would suggest that, due to context-dependent differential information sampling, it is the former (Vlaev, 2018). Thus, we have focused in this paper on experimentally elicited measures, as there at

least individuals largely make incentive-compatible (“real”) as opposed to hypothetical choices and probabilities are explicit, and therefore truly in the domain of risky decision making. The complexity of the matter at hand, however, still does not diminish. Researchers have shown that individuals adopt different strategies depending on the structure of the experiment (Pedroni et al, 2017). More damningly, even after differences in the structure of the experiment were controlled for, Pedroni et al were unable to elicit a stable measurement of risk preference, suggesting that an individual’s experimentally induced risk preference is likely constructed in the moment, multi-dimensional, and the product of multiple cognitive processes interacting. Overall, these variations appear to be largely a function of context, experimental structure, and the interaction of variable processes: very much in line with reasons to consider frameworks that explicitly account for contextual variability such as Vlaev’s local rationality as both plausible and appropriate.

Delineating Temporal Discounting

We note that less quantitative research has been conducted on the variability of measures relating to an individual’s temporal discounting, relative to risk.

Measurements. We can leverage the same overarching taxonomy to categorize intertemporal choice tasks as with risky decision making. In a typical behavioral-economic intertemporal choice task, individuals will choose between a smaller sooner “SS” option or a larger later “LL” option. Thus, experiments can be SDSC (e.g. titration methods when options presented depend on previous choice to arrive at ostensibly more precise estimates as in Solway et al, 2017), SIOC (e.g. options presented are independent of choice but have some structure, e.g. increasing LL option by \$5 each trial as in Steinglass et al, 2015) or SISC (most common: no built in cross-trial relationship e.g. Hunter et al, 2018). Unlike with risky decision making, intertemporal choice experiments about money are usually displayed only numerically – with monetary values and delay listed (e.g. Hunter et al, 2018). This ceases to strictly be the case when individuals make choices about non-monetary rewards like food or alcohol, where both pictorial representations and the physical objects they are choosing between can be presented (e.g. Addessi et al, 2014). Some experiments also vary choice domain (gain or loss, often in conjunction with gain/loss in risk, e.g. Estle et al, 2006);

incentivization and the amount of time an individual has to respond (Scherbaum et al, 2012). As there is no immediate uncertainty involved in pure intertemporal choice, feedback via choice realization as studied in risky decision making is largely inconsequential. Similarly, while some authors have considered the description-experience gap in intertemporal choice, this is typically only examined in the relatively uncommon context of probabilistic rewards (Dai et al, 2019): that is when either immediate and/or delayed rewards are themselves offered probabilistically. This makes sense as there is no ostensible learning or underlying uncertainty to be reduced besides inherent temporal uncertainty which is both outside the decision-maker's control and unable to be experienced – and thus seemingly resolved – till that moment in time. This is the case unless, for example, the experiment is situated in a virtual world where the experimenter is imperator and can control time.

Just as with risk preference, many temporal discounting elicitation methods exist in the literature: experiments, surveys (e.g. Barratt Impulsiveness Scale, Barratt & Patton, 1983), and frequency measures. *Variability.* Measurements of temporal discounting or impulsivity show varying degrees of temporal stability and predictive power. They are demonstrated to vary with affect (increases with sadness and reduces with gratitude: Lempert & Phelps, 2016); age (older adults are more patient: Green et al, 1994); attentional and framing manipulations (increases with focus on delay and decreases with focus on magnitude: Leiser et al, 2008, Lempert & Phelps, 2016); pathology (patients with substance use disorders are more impulsive: MacKillop et al, 2011); prospection (decreases with emphasis on future concreteness: Lempert & Phelps, 2016) and sex (women discount more steeply than men in the lab: Weafer & de Wit, 2014). Again, however, there is much extant literature suggesting a lack of systematic relationship (e.g. no conclusive direction one way or the other for sex differences: Cross et al, 2011). More concretely, human discounting, like risk preference, is sensitive to domain and circumstance – not only do people differentially discount across goods and money, they allocate fixed resources (money) on these goods depending on their current financial situation (Ubfal, 2016). In this study conducted in rural Uganda with non-hypothetical rewards, the less income an individual had, the more money they were willing to spend on items they discount highly. Similarly, individuals with gambling use disorder discounted money more highly when in

a gambling context, as opposed to a non-gambling context (Peters & Büchel, 2011). This is putative evidence for the influence of personal goals and contexts (income in the Ubfal example; physical location in the Peters & Büchel example) on choice behavior, something only speculated about earlier in this paper. Further, the test-retest reliability in measures of discount factor could also partially be dependent on reinstating the same context in which initial measurements were made: in a study, 5 week test-retest reliability of k was 0.77 (95% CI: 0.67-0.85, $n = 81$), 1 year was 0.71 (0.5-0.84, $n = 37$), and 57-weeks was 0.63 (0.41-0.77, $n = 46$) when subjects made choices between which delayed reward they preferred (Kirby, 2009). It would be interesting and might lead to better correspondence with real-world behavior to consider a within-subject design where such data were collected across multiple different contexts (e.g. in Ecological Momentary Assessment style experiments which we are sure must currently be in progress).

Unsurprisingly, inference on how an individual discounts value over time has also shown to be sensitive to the structure of the experiment. For example, Lempert and colleagues inferred different discount factors for subjects depending on how they manipulated stimuli in the experiment: people discounted significantly more steeply when there was greater variation in the delayed reward relative to the immediate reward but that the rank ordering of discount factors remained consistent regardless of experiment structure (Lempert et al, 2015). The researchers also had subjects complete various surveys measuring impulsivity / related factors and, like the Frey group, found significant correlations in temporal discounting measures derived from surveys but not between surveys and experiments. Other researchers have found similar evidence across multiple clinical populations. For example, Ledgerwood and group found these patterns held in control subjects and pathological gamblers with and without a history of substance use disorder (Ledgerwood et al, 2009). They further found that pathological gamblers were generally more impulsive regardless of substance use history, but that the gamblers with a history of substance use were more risk tolerant. This is just one simple example to demonstrate simultaneously the clinical significance of these economic constructs and how they may (or may not) vary across populations. That the overall relationship between methods of temporal-discounting elicitation seems to hold despite pathology suggests that, unless

we assume that this variability is irreducible, there may be other factors – cognitive or otherwise – that may not be considered carefully enough by the field.

Overall, we see remarkable correspondence in research studying the elicitation and sensitivity of both risk preference and temporal discounting. Regardless of the construct, there is much variability in the conceptualization and inference of parameter values. We further see that this variability tends to be greater in experimentally induced parameter inference. We next consider other observations of variability in choice behavior in experiments.

Variability in experiments: measurements

The idea that choice behavior and decision-making strategies in behavioral-economic experiments may change over the course of an experiment is not novel and has been explored for decades (Slovic, 1995; Vlaev 2018). When characterizing variability in behavior, we can, in general, consider macro (e.g. domain differences) and micro levels. Framing effects, like those described by Prospect Theory where people use a single reference point to guide behavior in the gain domain compared to loss, are examples of macro influences in that they demarcate domains (Kahneman & Tversky, 1979). Such reference points are traditionally assumed to be fixed – these frames may impact how an individual behaves on aggregate and on a given trial (but see Koop & Johnson, 2012 for empirical observations suggesting multiple reference points). Extant literature has shown that context can also exert a macro level influence on behavior and inferred parameters (Peters & Büchel, 2011). Context – depending on how it is defined – however, is particularly precarious and can also influence decision-making at micro-levels (e.g. Lempert et al, 2015) and in non-human primates (Zimmermann et al, 2018).

We can, then, decompose micro effects into trial, within-trial, and between-trial levels. Trial level measurements include the gold standard but highly variable choice behavior, and reaction time (though note that models of reaction time themselves are within trial as they seek to describe dynamics over the course of the trial itself). Choice behavior is generally modeled in accordance with SEUT or CPT maximization as described in the first section of this paper in conjunction with a choice rule.

Within-trial measurements generally include response time models and process-tracing methods such as mouse and eye tracking (see Schulte-Mecklenbeck et al, 2017 for more). Response times are widely modeled using a sequential sampling framework which assumes that we accumulate information in favor of/against the options presented to us in a noisy manner until we have accrued enough to make a choice (or never accrue enough to ever make a choice) (Forstmann et al, 2016). One of the most widely used frameworks, the Drift Diffusion Model, breaks down the accumulation process into four parameters in two alternative forced choice tasks: bias (predisposition towards Option A or B), drift rate (the rate at which evidence is accumulated), threshold (the amount of information needed to make a choice), and non-decision time (generally considered to be irrelevant to the decision process) (Ratcliff, 1978). These psychologically interpretable parameters that model components of deliberation have shown to correlate with discount factor (Hunter et al, 2018; Konovalov & Krajbich, 2019), risk tolerance (Konovalov & Krajbich, 2019), and preference more broadly (Konovalov & Krajbich, 2019).

Process-tracing methods, then, are direct measurements of the dynamics in decision making, capturing the online formation or reversal of preference (Koop & Johnson, 2013) and the mitigation of conflict as individuals choose between two options (Stillman et al, 2020). Scientists can measure and model these dynamics in decision making by examining the path subjects take via their computer mouse: a direct and swift movement from trial start to the option selected suggests decisive choice, whereas more winding trajectories could indicate decision difficulty (conflict) and even preference reversal. Such measurements further allow arbitration between different theories of preference formation. We omit discussion of eye tracking and value-based decision making due to space constraints (see Orquin & Loose, 2013 for a review) and instead focus on mouse tracking. For example, Koop and Johnson demonstrated that preference reversals during risky choice inferred via mouse tracking were inconsistent with heuristic decision strategies like “take the best” (TTB) which posits that individuals focus on a particular dimension and select the choice that ranks highest on that dimension (Koop & Johnson, 2013). TTB’s incompatibility with preference reversals were demonstrated by the degree with which mouse trajectories deviated from relative linearity (i.e. moving the mouse directly to the object of choosing). Similarly, Stillman and group found in

a risky decision-making task that the more similar the subjective values of choice options, the less direct and more conflicted the subjects' trajectories were despite controlling for response time greater the conflict (Stillman et al, 2020). The authors argued further that mouse trajectories could correlate with an individual's risk tolerance: an individual who follows a direct trajectory to the gamble as opposed to the certain option is likely more risk tolerant than someone who takes a meandering path. Stunningly, the authors found that decision conflict on single trials correlated strongly with risk tolerance, inferred in accordance with the Prospect Theory framework, and predicted behavior on the subsequent decision (Stillman et al, 2020). The authors argue that mouse-tracking dependent inference outperforms traditional behavioral measures of choice behavior and reaction time analyses as mouse-tracking might be more robust to other factors known to affect response time and choice behavior (e.g. non-decision time). Similarly, scientists have correlated decision strategy as inferred through mouse tracking dynamics with discount factor in intertemporal choice (Reeck et al, 2017). Such analyses suggest a promising avenue to understand more about locally rational, online decision making, and especially the role of similarity between options presented on a given trial.

Trial-Trial Temporal Dependence

Any discussion on human behavior and rationality would be incomplete without a brief further comment on capacity constraints and adaptive behavior. A key signifier of 'intelligence' is the ability to navigate complicated environments. Animal – and artificial – behavior is however hardware constrained: there are limits to the ability and flexibility that organisms and algorithms can demonstrate. Many theories of how the human brain evolved to be able to maneuver such a complicated world given limited resources revolve around the idea of adapting to or leveraging (stationary) statistical information in the environment (Anderson, 1990). In perceptual neuroscience, this is referred to as the Efficient Coding Hypothesis where limited probabilistic neuronal representations maximize information and minimize redundancy (Barlow, 1961) in a context-sensitive way (Schwartz et al, 2007). Such adaptive sequential sensitivity has been demonstrated in lower-level cognition (Simoncelli & Olshausen, 2001) and, more recently, in non-human primate economic decision making (Zimmermann et al, 2018). The results from the Zimmermann paper are

in particular valuable because the authors demonstrate the first evidence of [the necessity of] trial-trial temporal dependencies in canonical neuronal computations during economic choice. That is, not only does behavior in economic choice change as a function of variability in rewards (a prediction of the Efficient Coding Hypothesis), but models of neuronal computations that are consistent with efficient coding – typically specified at the intra-trial level – can *only* describe behavior well if the temporal order and (local) contexts of the experiment are preserved and explicitly accounted for. Taken together with theoretical and empirical neuroscientific research on sequential sampling in the brain (e.g. Gold & Shadlen 2007), this suggests that in the realm of rationality – resource or economic – context is king, and hence lends further credence to frameworks like Vlaev’s local rationality. Indeed, sequential sampling models like the Drift Diffusion Model have been monumentally successful in describing behavior alone (Forstmann et al, 2015).

There is, therefore, a strong intuition as to the normative reliance on recent history during experience in the moment – be it simply perceiving stimuli or during the decision process and subsequent choice itself. The mechanisms through which this might manifest are still fundamental open questions in the field, though there is general speculation on the (often complementary) roles of attention and working memory in propagating this temporal continuity (Kiyonaga et al, 2017). While this is usually considered to be adaptive (Kiyonaga et al, 2017), we highlight two cases in which reliance on recent history can prove to be problematic or task incongruent: environments without sequential dependencies and clinical pathology.

Much research has examined how reliance on the past can cause problems in lower-order cognition due to task-irrelevancy. Some of the earliest evidence of this comes from the absolute identification literature in the 1950s and onwards: where individuals were demonstrated to treat independently generated stimuli (i.e. presented a sequence of stimuli that were not related by time, like in SISC tasks) as if they were actually related (e.g. Verplanck & Blough, 1958; Lockhead & King, 1983, Stewart et al, 2005). For example, when people were asked to make judgements about line lengths or tone frequency, experimenters found robust evidence of transient framing effects: the lines or tones they had seen immediately (1 – 4 lags) before influenced their judgements on the current stimulus shown (Stewart et al, 2005). Interestingly, some experiments have shown different effects as a function of lag: more recent stimuli tend to produce an

attractor-style effect, while more distantly observed stimuli produce a contrast effect (Stewart et al, 2005). Researchers in visual perception have termed this effect, ostensibly distinct from priming, hysteresis, statistical artefacts, and learning, as *serial dependence* (Fischer & Whitney, 2014). This effect, consistent with the Efficient Coding Hypothesis, is also thought to be adaptive despite any inferential obstruction it may cause in such randomized tasks. We note that serial dependence has important consequences in real-world contexts too, and not just as a potential ‘contaminator’ of psychological inference. Recently, work from David Fischer’s group showed that radiologists demonstrated serial dependence while making medical judgements about simulated patient lesions (Manassi et al, 2021). More broadly, Fischer and Whitney suggest that serial dependence is characterized along three dimensions: similarity (only present when stimuli have similar features), temporality (decays over time), and spatiality (strongest when stimuli presented in the same location) (Fischer & Whitney, 2014). The authors also identify attention as a fundamental player (Fischer & Whitney, 2014). It is still, however, an open question as to whether this type of between-trial effect extends to higher-order (behavioral-economic) decision making.

In SISC (randomized) tasks, the standard experimental structure in intertemporal choice and risky decision making, stimuli are expectation controlled and thus presented in a randomized fashion. That is, successive stimuli will possess varying degrees of similarity to each other. For example, on trial $t - 1$ an individual chooses between \$5 today and \$45 in 80 days and on trial t chooses between \$4 today and \$48 in 70 days. Here the immediate and delayed rewards are similar in value, as is the delay of the rewards. Further, as choice options are often displayed in similar spatial locations (though we note that there is usually randomization at least in terms of the side of the screen – left or right – each option is presented), and decisions are made in a sequence, the criteria for plausible serial dependence according to Fischer and Whitney appear to be met. It is entirely conceivable and ostensibly efficient for computations made during trial $t - 1$ to be (partially) cached and reused or referenced on trial t as a function of similarity (Dasgupta et al, 2018), amongst other things, thus affecting response times and possibly choice behavior. Such influences, however transient they may be, may provide us information as to the cognitive health of an individual (see below) and may aggregate to the point of affecting inference on our parameters of interest,

especially if they are not accounted for in statistical analyses. Indeed, some of the concerns raised in the *Delineating* sections earlier with regards to noisiness in experimentally-induced parameter inference, could be due in part to such spillover. These spillover effects may also be consistent with cognitive theories of intertemporal and risky choice. As intertemporal choice involves uncertainty about the future, Peters and Büchel (amongst others) suggest that people's choices are guided by the deliberative process of prospection – they imagine what their future may look like some d days out and use the outcome of that simulation to guide their choice (Peters & Büchel, 2011). Recent work has also shown a relationship between how model-based an individual is and the way they discount the future: people who spend more time considering future rewards in temporal-discounting tasks are also more likely to plan ahead in sequential reinforcement learning tasks (Hunter et al, 2018). Further, scientists have hypothesized that the manner in which people choose also changes as a function of delay: Construal Level Theory posits that representations of the future are more abstract (e.g. lower statistical precision) than representations of the present (Lieser & Hadar, 2008) and that people tend to consider more “primary” attributes (e.g. healthiness, “should” behaviors) when thinking of the future and more “secondary” attributes (e.g. tastiness, “want” behaviors) for the present (Rogers & Bazerman, 2008). We can therefore infer that thinking carefully about anything – in this case the future – can be resource intensive. Thus, in the example above, the individual has already imagined what their life might look like 80 days into the future on the previous trial. Barring some specific event that they expect to meaningfully shape their experience within the 10-day difference, it is likely that their future 70 days out will be similar and thus they could avoid computational redundancy by reusing (part of) the simulation generated on the previous trial. Indeed, if representations of the future are in actuality more uncertain and requiring the recruitment of higher-order cognitive processes, there is even more reason to support the reuse of previous computations to guide current inference and choice to minimize computationally expensive operations. While decision making under risk may not involve the simulating the future, individuals still need to resolve the immediate uncertainty and complex choice options presented in the form of probabilistic gambles to guide their choice. Thus, computations incurred over the course of the experiment may likewise be carried over from trial to trial, also possibly as a function of (dis)similarity.

Our recent work introduces a novel statistical framework that suggests choice behavior, response times, and risk tolerance/discount factors themselves are indeed influenced by recent history as defined by previous stimuli and choices made (Banavar & Bornstein, *PsyArxiv*). We term this dependence *computational perseveration* to distinguish its higher-order nature (involving complex mental calculations) from serial dependence. We find specific effects of computational perseveration in choice behavior, while reaction time parameters showed more widespread sensitivity. However, our results suggest further complexity in the nature of this higher-order serial dependence as we also found evidence for a contrast-like effect: in the risky decision-making task, choice behavior was influenced by previous stimuli when the previous choice was easy, and the current was difficult. This is the opposite of what would be expected given Fischer & Whitney's criteria (Fischer & Whitney, 2014). Critically, our analyses have shown that the majority of subjects in an Intertemporal Choice task and in a Risk/Ambiguity gambling task show evidence of computational perseveration. Further, we demonstrate that for several subjects, sequential-effect-adjusted parameters for risk and ambiguity tolerance change sign, and therefore, psychological interpretation. For example, someone who was previously identified as 'ambiguity seeking' based on their non-sequential-effect adjusted ambiguity tolerance parameter would now be identified as 'ambiguity averse.' We argue therefore that while computational perseveration may not be the sole source of variability in risk tolerance/discount factor inference in experiments, there is theoretical and empirical impetus for us to consider explicitly the influence of temporal context in how we define, measure, and *infer* these constructs. We finally note that computational perseveration is likely present in SDSC and SIOC tasks too, but due to the potentially confounding nature of structured experiments and learning, we omit further consideration of this topic in this paper.

We believe that this work has deep theoretical and empirical implications. Our analyses suggest that these sequential effects are not noisy artefacts but are instead the consequence of a systematic influence of trial properties on components of the decision process. This suggests a potential need for the theoretical reconceptualization of *experimentally-inferred* parameters as *explicitly dynamic* and sensitive to (highly) local contexts and not *exclusively* a static and psychologically interpretable end (Banavar & Bornstein,

PsyArxiv). Our method also allows scientists to analyze a novel dimension of information about the decision maker (i.e. degree of trial-trial dependencies) without having to collect any new measures, as both choice behavior and response times are standardly recorded in experiments. This additional information could have use beyond the purely methodological – it could result in meaningful cognitive and clinical implications.

To underscore the idea that short-term temporal dependencies provide cognitively meaningful information and other research directions beyond parameter calibration (as it may be tempting to infer from the previous paragraph), we consider in brief a complementary, yet distinct, line of work in clinical psychology and neurology. Decades of evidence in these fields has shown a differential reliance on recent history as a function of aging and neurodegenerative pathology (Sandson & Albert, 1984; Goldberg, 1986; van Patten et al, 2015). Here the abnormal, often over- and task-incongruent reliance on the past, relative to healthy individuals, is termed perseveration. In particular, there exists a three-dimensional hierarchy of perseveration with primary dimensions of content, disorder, and temporal profile. Content references the material itself that is repeated (ranging from lower-order motor to higher-order semantic/verbal repetitions); disorder references the various ways in which outcome measures might differentially relate to the clinical progression of neural degeneration (e.g. frontal lobe vs basal ganglia damage) and temporal profile, which delineates the varying timeframes along which perseveration can manifest (e.g. perseverate information from seconds ago, minutes ago, or even tasks ago) (Sandson & Albert, 1984; Goldberg, 1986; Serpell et al, 2009; van Patten et al, 2015). Like with serial dependence in visual perception, the upper limits of the content hierarchy are unknown, and a future line of research examining the presence or absence of computational perseveration – the degree to which there is dependence on the recent past – in aging and disease during complex decision making may lead to a novel marker of cognitive decline.

Conclusion

Preferences are by definition subjective. Decades of research into risk preference and temporal discounting have conclusively shown that these concepts, however they may be defined, instantiated, or measured, are variable. One extensive form of variability comes from the multiple well-established methods to elicit these

measurements – often either in an experiment or by completing surveys. Further, individuals (and subsequently inferred parameters) demonstrate sensitivity to domains, context (recent history in both choice and stimulus, environmental uncertainty), and adaptation (e.g. shifting reference points, preference reversals). These sensitivities have been demonstrated in healthy individuals and clinical populations, with often meaningful differences between groups. In the growing field of computational psychiatry there is much research focused on linking measures of risk preference and temporal discounting to maladaptive behavior. While there has been much success on this front, understanding and appropriately characterizing these concepts in health and disease is critical. In this paper, we have reviewed some of the different ways in which these concepts have been characterized and operationalized and have proposed another source of variability that we believe deserves further scrutiny: the explicit influence of recent history on choice behavior, response times, and subsequently inferred values.

We suggest that such trial-level sequential influences are adaptive and consistent with ideas of contextual or local rationality. Ample evidence in the psychophysics and perceptual decision-making literature (amongst others) demonstrates that even when all pains are taken to minimize sequential dependencies within an experiment, the seriality of our temporal experience [and neural processing] plays a profound, arguably causal and adaptive role in shaping behavior. People’s fundamental conceptualizations about parameters and constructs are largely shaped by the functional forms and methods used to describe and infer them – compare classical Bernoulli Utility to Random Utility models, or evidence accumulation models with and without noisy accumulation of evidence, for example. By incorporating trial order (and recent history more generally) into the modeling of risk preferences and discount factors themselves, we hope that the field will move more concretely towards embracing these concepts as inherently, and therefore necessarily, contextual. This could lead towards better reconciling myriad behavioral observations and moving towards a more veridical notion of human rationality.

Open practices statement: The data and materials used in the paper are available from the corresponding author on reasonable request. No experiments were preregistered.

References

- Addressi, E., Bellagamba, F., Delfino, A., De Petrillo, F., Focaroli, V., Macchitella, L., Maggiorelli, V., Pace, B., Pecora, G., Rossi, S., Sbaffi, A., Tasselli, M.I., & Paglieri, F. (2014). Waiting by mistake: Symbolic representation of rewards modulates intertemporal choice in capuchin monkeys, preschool children and adult humans. *Cognition*. 130(3), 428-441. <https://doi.org/10.1016/j.cognition.2013.11.019>.
- Allais M. (1990) Allais Paradox. In: Eatwell J., Milgate M., Newman P. (eds) *Utility and Probability*. The New Palgrave. Palgrave Macmillan, London. https://doi.org/10.1007/978-1-349-20568-4_2
- Anderson, J. R. (1990). *Studies in Cognition. The adaptive character of thought*. Lawrence Erlbaum Associates, Inc.
- Andreoni, A., Callen, M., Hussain, K., Khan, M.Y., & Sprenger, C. (2016). Using preference estimates to customize incentives: an application to polio vaccination drives in Pakistan. *Journal of the European Economic Association*.
- Banavar, N. V., & Bornstein, A. (2023, February 27). Independent, not irrelevant: Trial order causes systematic misestimation of economic choice traits. *PsyArxiv*. <https://doi.org/10.31234/osf.io/a8gz3>
- Banavar, N.V., Lee, M.D., & Bornstein, A.M. (2021). Sequential effects in non-sequential tasks. *Proceedings of the 19th International Conference on Cognitive Modeling*.
- Barlow, H.B. (1961). Possible Principles Underlying the Transformation of Sensory Messages. *Sensory Communication*. 217-234
- Barratt, E.S., & Patton, J.H. (1983). Impulsivity: cognitive, behavioral, and psychophysiological correlates. In: Zuckerman, H (ed) *Biological bases of sensation seeking, impulsivity, and anxiety*. Erlbaum, Hillsdale, NJ

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16(3), 215–233.

<https://doi.org/10.1002/bdm.443>

Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76, 412–427.

<https://doi.org/10.1016/j.neuroimage.2013.02.063>

Bickel, W. K., Koffarnus, M. N., Moody, L., & Wilson, A. G. (2014). The behavioral- and neuro-economic process of temporal discounting: A candidate behavioral marker of addiction. *Neuropharmacology*, 76, 518–527.

<https://doi.org/10.1016/j.neuropharm.2013.06.013>

Brooks, H. R., & Sokol-Hessner, P. (2020). Quantifying the immediate computational effects of preceding outcomes on subsequent risky choices. *Scientific Reports*, 10(1), 9878.

<https://doi.org/10.1038/s41598-020-66502-y>

Crosan, R., & Gneezy, U. (2009). Gender Differences in Preferences. *Journal of Economic Literature*, 47(2), 448–474.

<https://doi.org/10.1257/jel.47.2.448>

Cross, C. P., Copping, L. T., & Campbell, A. (2011). Sex differences in impulsivity: A meta-analysis. *Psychological Bulletin*, 137(1), 97–130.

<https://doi.org/10.1037/a0021591>

Dai J, Pachur T, Pleskac TJ, & Hertwig R. (2019). What the Future Holds and When: A Description–Experience Gap in Intertemporal Choice. *Psychological Science*, 30(8), 1218–1233.

doi:10.1177/0956797619858969

Dasgupta, I., Schulz, E., Goodman, N.D., & Gershman, S.J. (2018). Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition*, 178, 67–81.

Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, 75(4), 643-669.

Estle, S.J., Green, L., Myerson, J. & Holt, D.D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. *Memory & Cognition*, 34, 914–928.
<https://doi.org/10.3758/BF03193437>

Evenden, J. L. (1999). Varieties of impulsivity. *Psychopharmacology*, 146(4), 348–361.
<https://doi.org/10.1007/PL00005481>

Gilboa, I., & Schmeidler, D. (1989). "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, Elsevier, vol. 18(2), pages 141-153, April

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, 17(5), 738–743. <https://doi.org/10.1038/nn.3689>

Fishburn, P., & Rubinstein, A. (1982). Time Preference. *International Economic Review*, 23(3), 677-694.
doi: 10.2307/2526382

Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, 3(10), e1701381.
<https://doi.org/10.1126/sciadv.1701381>

Frey, R., Richter, D., Schupp, J., Hertwig, R., & Mata, R. (2020). Identifying robust correlates of risk preference: A systematic approach using specification curve analysis. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspp0000287>

Fründ, I., Wichmann, F. A., & Macke, J. H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision*, 14(7), 9–9. <https://doi.org/10.1167/14.7.9>

Forstmann, B.U., Ratcliff, R., & Wagenmakers, E.J. (2016). Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions. *Annual Review of Psychology*, 67, 641-666

Gold, J. I., & Shadlen, M. N. The Neural Basis of Decision Making. *Annual Review of Neuroscience*, 30(1), 535-574.

Goldberg, E. (1986). Varieties of Perseveration: A Comparison of Two Taxonomies. *Journal of Clinical and Experimental Neuropsychology*, 8(6), 710-726.

Green, L., Fry, A. F., & Myerson, J. (1994). Discounting of Delayed Rewards: A Life-Span Comparison. *Psychological Science*, 5(1), 33–36. <https://doi.org/10.1111/j.1467-9280.1994.tb00610.x>

Guan, M., Stokes, R., Vandekerckhove, J., & Lee, M. D. (2020). A Cognitive Modeling Analysis of Risk in Sequential Choice Tasks. *Judgment and Decision Making*, 24.

Hayden, B. Y., & Platt, M. L. (2007). Temporal Discounting Predicts Risk Sensitivity in Rhesus Macaques. *Current Biology*, 17(1), 49–53. <https://doi.org/10.1016/j.cub.2006.10.055>

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517–523. <https://doi.org/10.1016/j.tics.2009.09.004>

Huettel, S. A., Stowe, C. J., Gordon, E. M., Warner, B. T., & Platt, M. L. (2006). Neural Signatures of Economic Preferences for Risk and Ambiguity. *Neuron*, 49(5), 765–775. <https://doi.org/10.1016/j.neuron.2006.01.024>

Hunter, D. R. (2002). Risk Perception and Risk Tolerance In Aircraft Pilots: (646002007-001) [Data set]. American Psychological Association. <https://doi.org/10.1037/e646002007-001>

Hunter, L. E., Bornstein, A. M., & Hartley, C. A. (2018). A common deliberative process underlies model-based planning and patient intertemporal choice [Preprint]. *Animal Behavior and Cognition*. <https://doi.org/10.1101/499707>

Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, 10(12), 1625–1633. <https://doi.org/10.1038/nn2007>

Kahneman, D., & Thaler, R. (2006). "Anomalies: Utility Maximization and Experienced Utility." *Journal of Economic Perspectives*, 20 (1): 221-234.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, Vol. 47, No. 2. (Mar., 1979), pp. 263-292.

Kirby, K.N. (2009). One-year temporal stability of delay-discount rates. *Psychonomic Bulletin & Review*, 16, 457-462.

Kiyonaga, A., Scimeca, J.M., Bliss, D.P., & Whitney, D. (2017). Serial Dependence across Perception, Attention, and Memory. *Trends in Cognitive Sciences*, 21(7), 493-497

Konova, A. B., Lopez-Guzman, S., Urmanche, A., Ross, S., Louie, K., Rotrosen, J., & Glimcher, P. W. (2020). Computational Markers of Risky Decision-making for Identification of Temporal Windows of Vulnerability to Opioid Use in a Real-world Clinical Setting. *JAMA Psychiatry*, 77(4), 368. <https://doi.org/10.1001/jamapsychiatry.2019.4013>

Kononov, A., & Krajbich, I. (2019). Revealed strength of preference: Inference from response times. *Judgement and Decision Making* 14(4), 381-394.

Koop, G. J., & Johnson, J. G. (2012). The use of multiple reference points in risky decision making. *Journal of Behavioral Decision Making*, 25(1), 49–62. <https://doi.org/10.1002/bdm.713>

Koop, G. J., & Johnson, J. G. (2013). The response dynamics of preferential choice. *Cognitive Psychology*, 67(4), 151–185. <https://doi.org/10.1016/j.cogpsych.2013.09.001>

Kugler, T., Connolly, T., & Ordóñez, L. D. (2012). Emotion, Decision, and Risk: Betting on Gambles versus Betting on People. *Journal of Behavioral Decision Making*, 25(2), 123–134. <https://doi.org/10.1002/bdm.724>

Lee, M.D. (2018). Bayesian Methods in Cognitive Modeling. In Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience, J.T. Wixted (Ed.). <https://doi.org/10.1002/9781119170174.epcn502>

Ledgerwood, D.M., Alessi, S. M., Phoenix, N., Petry, N.M. (2009). Behavioral assessment of impulsivity in pathological gamblers with and without substance use disorder histories versus healthy controls. *Drug and Alcohol Dependence*, 105. 89 – 96

Leiser, D., Azar, O. H., & Hadar, L. (2008). Psychological construal of economic behavior. *Journal of Economic Psychology*, 29(5), 762–776. <https://doi.org/10.1016/j.joep.2008.08.002>

Lempert, K. M., Glimcher, P. W., & Phelps, E. A. (2015). Emotional arousal and discount rate in intertemporal choice are reference-dependent. *Journal of Experimental Psychology. General*, 144(2), 366–373. <https://doi.org/10.1037/xge0000047>

Lempert, K. M., & Phelps, E. A. (2016). The Malleability of Intertemporal Choice. *Trends in Cognitive Sciences*, 20(1), 64–74. <https://doi.org/10.1016/j.tics.2015.09.005>

Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural Representation of Subjective Value Under Risk and Ambiguity. *Journal of Neurophysiology*, 103(2), 1036–1047. <https://doi.org/10.1152/jn.00853.2009>

Liberman, N., & Trope, Y. (2003). Construal level theory of intertemporal judgment and decision. In G. Loewenstein, D. Read, & R. Baumeister (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (p. 245–276). Russell Sage Foundation.

Loewenstein, G., O'Donoghue, T., & Bhatia, S. (2015). Modeling the interplay between affect and deliberation. *Decision*, 2(2), 55–81. <https://doi.org/10.1037/dec000029>

Lopez-Guzman, S., Konova, A. B., Louie, K., & Glimcher, P. W. (2018). Risk preferences impose a hidden distortion on measures of choice impulsivity. *PLOS ONE*, 13(1), e0191357. <https://doi.org/10.1371/journal.pone.0191357>

MacKillop, J., Amlung, M. T., Few, L. R., Ray, L. A., Sweet, L.H., & Munafò, M. R. (2011). Delayed reward discounting and addictive behaviors: a meta-analysis. *Psychopharmacology* 216, 305 – 321. <https://doi.org/10.1007/s00213-011-2229-0>

Manassi, M., Ghirardo, C., Canas-Bajo, T., Ren, Zhihang, Prinzmetal, W., & Whitney, D. (2021). Serial dependence in the perceptual judgements of radiologists. *Cognitive Research: Principles and Implications*, 6

Orquin, J. L., & Loose, S. M. (2013). Attention and choice: A review on eye movements in decision making. *Acta psychologica*, 144(1), 190-206.

Pedroni, A., Frey, R., Bruhin, A., Dutilh, G., Hertwig, R., & Rieskamp, J. (2017). The risk elicitation puzzle. *Nature Human Behaviour*, 1(11), 803–809. <https://doi.org/10.1038/s41562-017-0219-x>

Pekkala, S., Albert, M.L., Spiro III, A., & Erkinjuntti, T. (2008). Perseveration in Alzheimer's disease. *Dementia and Geriatric Cognitive Disorders*, 25, 109-114.

Peters, J., & Büchel, C. (2011). The neural mechanisms of inter-temporal decision-making: Understanding variability. *Trends in Cognitive Sciences*, 15(5), 227–239. <https://doi.org/10.1016/j.tics.2011.03.002>

Pulcu, E., Trotter, P. D., Thomas, E. J., McFarquhar, M., Juhasz, G., Sahakian, B. J., Deakin, J. F. W., Zahn, R., Anderson, I. M., & Elliott, R. (2014). Temporal discounting in major depressive disorder. *Psychological Medicine*, 44(9), 1825–1834. <https://doi.org/10.1017/S0033291713002584>

Radulescu, A., Holmes, K., & Niv, Y. (2020). On the convergent validity of risk sensitivity measures. *PsyArXiv*. <https://doi.org/10.31234/osf.io/qdhx4>

Ratcliffe, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59-108.

Reeck, C., Wall, D., & Johnson, E.J. (2017). Search predicts and changes patience in intertemporal choice. *Proceedings of the National Academy of Sciences*, 114(45), 11890-11895.

Rogers, T., & Bazerman, M. H. (2008). Future lock-in: Future implementation increases selection of ‘should’ choices. *Organizational Behavior and Human Decision Processes*, 106(1), 1–20. <https://doi.org/10.1016/j.obhdp.2007.08.001>

Sandson, J. & Albert, M.L. (1984). Varieties of perseveration. *Neuropsychologia*, 22(6), 715-732.

Samuelson, P. (1937). A Note on Measurement of Utility. *The Review of Economic Studies*, 4(2), 155-161. Retrieved January 14, 2021, from <http://www.jstor.org/stable/2967612>

Savage, L. J. (1954). *The Foundations of Statistics*. Naval Research Logistics, 1: 236-236. John Wiley & Sons, Inc.

Scherbaum, S., Dshemuchadse, M., Goschke, T. (2012). Building a Bridge into the Future: Dynamic Connectionist Modeling as an Integrative Tool for Research on Intertemporal Choice. *Frontiers in Psychology*, 3.

Schulte-Mecklenbeck, M., Johnson, J. G., Böckenholt, U., Goldstein, D. G., Russo, J. E., Sullivan, N. J., & Willemsen, M. C. (2017). Process-Tracing Methods in Decision Making: On Growing Up in the 70s. *Current Directions in Psychological Science*, 26(5), 442–450. <https://doi.org/10.1177/0963721417708229>

Schwartz, O., Hzu, A., & Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience*, 8, 522-545

Serpell, L., Waller, G., Fearon, P., & Meyer, C. (2009). The Roles of Persistence and Perseveration in Psychopathology. *Behavior Therapy*, 40, 260-271.

Simon, H.A. (1990). Bounded Rationality. In *Utility and Probability*, eds. Eatwell, J., Milgate, M., & Newman, P. Palgrave Macmillan, London.

Simoncelli, E. P., & Olshausen, B.A. (2001). Natural Image Statistics and Neural Representation. *Annual Review of Neuroscience*, 24, 1193-1216

Slovic, P. (1995). The construction of preference. *American Psychologist*, 50(5), 364–371. <https://doi.org/10.1037/0003-066X.50.5.364>

Sokol-Hessner, P., Raio, C. M., Gottesman, S. P., Lackovic, S. F., & Phelps, E. A. (2016). Acute stress does not affect risky monetary decision-making. *Neurobiology of Stress*, 5, 19–25. <https://doi.org/10.1016/j.ynstr.2016.10.003>

Solway, A., Lohrenz, T., & Montague, P. R. (2017). Simulating future value in intertemporal choice. *Scientific Reports*, 7(1), 43119. <https://doi.org/10.1038/srep43119>

Steinglass, J. E., Lempert, K. M., Choo, T.-H., Kimeldorf, M. B., Wall, M., Walsh, B. T., Fyer, A. J., Schneier, F. R., & Simpson, H. B. (2017). Temporal discounting across three psychiatric disorders: Anorexia nervosa, obsessive compulsive disorder, and social anxiety disorder. *Depression and Anxiety*, 34(5), 463–470. <https://doi.org/10.1002/da.22586>

Stewart, N., Brown, G. D. A., & Chater, N. (2005). Absolute Identification by Relative Judgment. *Psychological Review*, 112(4), 881–911. <https://doi.org/10.1037/0033-295X.112.4.881>

Stillman, P. E., Krajbich, I., & Ferguson, M. J. (2020). Using dynamic monitoring of choices to predict and understand risk preferences. *Proceedings of the National Academy of Sciences*, 117(50), 31738–31747. <https://doi.org/10.1073/pnas.2010056117>

Thaler, R., 1981. Some empirical evidence on dynamic inconsistency. *Economics Letters*, 201–207

Thurstone., L. L. (1927) A law of comparative judgement. *Psychological Review*, 34(4):273–286.

Tversky A., & Kahneman D. (1989) Rational Choice and the Framing of Decisions. In: Karpak B., Zionts S. (eds) *Multiple Criteria Decision Making and Risk Analysis Using Microcomputers*. NATO ASI Series (Series F: Computer and Systems Sciences), vol 56. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-74919-3_4

Tymula, A., Glimcher, P. W., Levy, I., & Belmaker, L. A. R. (2012). Separating Risk and Ambiguity Preferences Across the Life Span: Novel Findings and Implications for Policy. 66.

Ubfal, D. (2016). How general are time preferences? Eliciting good-specific discount rates. *Journal of Development Economics*, 118, 150–170. <https://doi.org/10.1016/j.jdeveco.2015.07.007>

Vlaev, I. (2018). Local Choices: Rationality and the Contextuality of Decision-Making. *Brain Sciences*, 8(1), 8. <https://doi.org/10.3390/brainsci8010008>

Weafer, J., & de Wit, H. (2014). Sex-differences in impulsive action and impulsive choice. *Addictive Behaviors*, 39(11), 1573 – 1579.

Weber, E. U. (2004). Perception matters: Psychophysics for economists. In Carrillo, J., Brocas, I. (Eds.), *Psychology and economics* (pp. 165–176). Oxford, England: Oxford University Press.

Weber, E. U. (2010). Risk attitude and preference. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(1), 79–88. <https://doi.org/10.1002/wcs.5>

Weber, E. U., Blais, A.-R., & Betz, N. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, 15, 263-290

van Patten, R., Kaufman, D.A.S., Mitchell, S., Sachs, B., & Loring, D.W. (2015). Perseverative Error Subtypes in Patients with Alzheimer's Disease and Mild Cognitive Impairment. *Journal of Neurology and Psychology*. Special Issue.

Verplanck, W. S., & Blough, D. S. (1958). Randomized Stimuli and the Non-Independence of Successive Responses at the Visual Threshold. *The Journal of General Psychology*, 59(2), 263–272. <https://doi.org/10.1080/00221309.1958.9710195>

von Neumann, J. & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

Zimmermann, J., Glimcher, P.W., & Louie, K. (2018). Multiple timescales of normalized value coding underlie adaptive choice behavior. *Nature Communications*, 9. <https://doi.org/10.1038/s41467-018-05507->