

THE COGNITIVE TURN IN BEHAVIORAL ECONOMICS*

Benjamin Enke

July 10, 2024

Preliminary version – comments and additional references very welcome

Abstract

This review organizes the vibrant recent literature on the cognitive foundations of economic decision-making. At a basic level, this entire literature studies imperfections in cognitive information processing: the ways in which people attend to, remember, aggregate and trade off variables to make economic decisions. A main idea in this literature is that many ostensibly-distinct empirical regularities and anomalies reflect generic simplification strategies that people adopt to reduce information-processing demands. These cognitive strategies can be consolidated into five categories: (i) noisy approximations and resulting behavioral attenuation; (ii) comparative thinking; (iii) reducing cardinality by overweighting what's salient, gets cued in memory, or is deemed important; (iv) thinking in analogies and categories; and (v) devaluing or shying away from objects one cannot properly evaluate. Work on information processing has both reinvigorated the upstream exchange with cognitive psychology and has started to trickle down into applied fields such as finance, labor and development. I discuss open questions, emphasizing both the need for a unified model that brings together the different simplification strategies emphasized in the literature; and a greater focus on economic applications and multi-agent settings.

*I thank Egor Bronnikov, Philipp Schirmer and Anna Valyogos for excellent research assistance. I am especially grateful to Ryan Oprea and Andrei Shleifer for many discussions on the topic. For comments I also thank Mark Dean, Thomas Graeber, Alex Imas, George Loewenstein, Pietro Ortoleva, Josh Schwartzstein, Zachary Wojtowicz and Florian Zimmermann. Enke: Harvard University, Department of Economics, and NBER, enke@fas.harvard.edu.

Contents

1	The Cognitive Turn	1
2	Overview: Information Processing and Simplification	5
2.1	Policy Uncertainty	5
2.2	Information Processing	7
2.3	Simplification Strategies	10
3	Noisy Approximations	12
3.1	Behavioral Attenuation	13
3.1.1	Evidence	13
3.1.2	Models of Behavioral Attenuation	18
3.2	What Determines the Magnitude of Noise?	21
3.3	Asymmetric Problem Setups and Spurious Behavioral Motivations	23
4	Comparative Thinking	24
4.1	Comparison and Reference Point Effects in Choice	25
4.1.1	Pairwise Comparisons	25
4.1.2	Comparisons to Normatively Irrelevant Reference Quantities	25
4.2	Quality Assessments	28
5	Reducing Cardinality	29
5.1	Evidence on Incomplete Representations	30
5.2	Goal-Directed Attention and Memory	33
5.3	Stimulus-Driven Attention	35
5.3.1	Choice	36
5.3.2	Beliefs	38
5.4	Similarity-Based Recall	40
5.4.1	Beliefs	40
5.4.2	Choice	43
6	Categorizations and Analogies	43
7	Complexity Aversion	46
8	What Determines Information Processing Imperfections?	47
8.1	The Complexity of Procedures and Algorithms	47
8.2	Features of the Decision Maker	49

9	Revisiting Some Classics	49
9.1	What We (Probably) Understand	50
9.2	Where More Work is Needed	52
9.3	What Seems Largely Unrelated	53
10	Evidence From the Field	54
11	Outlook and Open Questions	56
11.1	Theory	56
11.2	Economic Applications	57
11.3	Experiments	58

1 The Cognitive Turn

Research in behavioral economics (aka Psychology and Economics) is currently taking a major “cognitive turn.” Much recent work focuses on empirically identifying and formally modeling the cognitive foundations of economic decision making, such as attention, memory, cognitive noise, and the ways in which they depend on a problem’s complexity. This review (i) discusses why the cognitive turn emerged and why it may be important; (ii) provides initial evidence for its occurrence; and (iii) synthesizes and organizes the rapidly-growing body of work on the topic that has accumulated over the last few years.

Why did the cognitive turn occur? Various objective metrics – awards, publications, conferences, dissertations and graduate courses – suggest that behavioral economics is a big success story. This success partly reflects the field’s ability to empirically identify and model a large collection of behavioral “anomalies”, many of which have been documented to affect consequential economic decisions. However, the success of behavioral economics at identifying ever more ways in which the neoclassical model fails also caused a well-understood disadvantage – a perceived *proliferation of concepts* (e.g., Fudenberg, 2006; Spiegler, 2019). For example, we often employ different approaches for understanding choice under risk, intertemporal choice and consumer choice, even when the empirical regularities and psychological principles are very similar to each other. Related is a recurring concern over stability. Much recent evidence highlights that various supposedly-fixed behavioral preferences and biases strongly vary with contextual features such as complexity or cues, raising the question how we might ultimately be able to predict when and where behavioral anomalies occur.

As a result of these and related questions, the broad view of the movement that I here label the *cognitive turn* is that – while we continue to explore field applications of by-now established concepts – it is worth investigating whether a smaller set of common cognitive principles can be identified to push the field even further.

Characteristics of the cognitive turn. While the cognitive turn has seen different approaches, it shares multiple common features. First, the field is often (though not always) more focused on explaining and unifying anomalies than on accumulating new deviations from neoclassical predictions per se. As a result, the interrelationships of biases have attracted growing attention (Dean and Ortoleva, 2019; Chapman et al., 2023a; Stango and Zinman, 2023; Enke et al., 2024a). A basic premise that underlies much of this work is that many behaviors reflect imperfections in basic *cognitive information processing* – attending to, remembering, aggregating and trading off variables to make economic decisions.

Second, the field generally aims to replace reduced-form notions of biases or “revealed non-standard preferences” with the underlying cognitive mechanisms. In this regard, an overly simplistic but perhaps helpful summary is that upon observing a non-standard choice

pattern, the knee-jerk reaction that motivates much research on cognitive foundations is to ask: “Which cognitive limitation generates this choice behavior, and how does it resemble anomalies we’ve seen in other domains?”, rather than, for example: “Which utility function rationalizes this choice?” As a result, the field often blurs the distinction between non-standard preferences and beliefs. Until recently, behavioral economists typically deployed two different toolkits to understand anomalies in beliefs and choice behavior, a main reason being that choice anomalies were believed to largely reflect non-standard utility functions. The cognitive turn has shifted attention to the hypothesis that the same basic mechanisms of processing information shape both beliefs and choices.

Intellectual origins and antecedents. While the cognitive turn accelerated only recently (the vast majority of the papers discussed below were written within the last ten years), its intellectual origins trace back to various earlier lines of work.¹ The most important very early antecedents in economics are the bounded rationality movement pioneered by Simon (1956, 1982) and the “system 1 vs. system 2” literature summarized by Kahneman (2011).² Simon’s work emphasized the importance of information processing early on, and contributions such as prospect theory – while ultimately relatively reduced-form in nature – were also full of appeals to cognitive foundations from perceptual and cognitive psychology.

Despite these important early streams of work, cognitive foundations received relatively little attention as behavioral economics came of age. Indeed, it is conceivable that the more reduced-form approach that dominated behavioral economics was impactful precisely because it ignored the nitty-gritty of information-processing imperfections, and was thus more workable in economic theory and practice (see Rabin, 2013, for a powerful discussion along these lines).

Still, the 1990s and early 2000s saw various early highly original contributions that – in combination with Simon, Kahneman and Tversky – paved the way for the research that started taking off in the mid-2010s. These early contributions include, for example, theoretical work on cue-based recall and decision-making (Laibson, 2001; Mullainathan, 2002; Bernheim and Rangel, 2004), research on analogical reasoning and categorization (Gilboa and Schmeidler, 1995; Jehiel, 2005; Mullainathan et al., 2008), work on rational inattention and noisy cognition (Sims, 1998, 2003), research on thinking styles (Frederick, 2005), a wave of work on information acquisition using process tracing techniques (John-

¹The cognitive turn and its antecedents in economics build on various literatures in psychology and neuroscience. This includes, for example, research on exemplar models and similarity (Tversky, 1977; Kahana, 2012) that forms the basis for much current work on memory. Similarly, much of the contemporary work on noisy cognition builds on research on noisy informing processing in psychology and neuroscience, including work on resource-rational cognition (Lieder and Griffiths, 2020), Bayesian models of cognitive noise (Oaksford and Chater, 2007; Gershman, 2021) and accounts of efficient coding (Barlow, 1961).

²In contrast to much of earlier work on bounded rationality (e.g., Conlisk, 1996; Harstad and Selten, 2013; Gigerenzer and Todd, 1999; Gigerenzer and Gaissmaier, 2011), recent work often directly models and studies mechanisms such as attention and memory, and has often maintained the assumption of utility-maximization.




Figure 1: Publication trends in behavioral economics in the profession’s “Top 5” journals over time, binned into three-year buckets. Forthcoming papers are included in 2023. Data collection as of January 2024.

son et al., 2002; Camerer et al., 2005), work that pioneered much subsequent research on choice uncertainty (Ariely et al., 2003), and research on narrow bracketing and diminishing sensitivity that attempted to tie together behavior across different domains (e.g., Prelec and Loewenstein, 1991; Read et al., 1999).

To assess publication trends more systematically, I classified all behavioral economics papers that were published in the profession’s “top 5” journals between 2000 and 2023 into six categories: bounded rationality in individual decision making; social preferences; risk and time preferences; belief-based utility (including image concerns and motivated reasoning); strategic reasoning; and others (e.g., gender). Figure 1 shows an overview of the publication trends. The left panel shows absolute counts per category and the right panel shares. The overall picture is clear: research on bounded rationality significantly increased starting in the mid-2010s, mostly at the expense of work on strategic thinking and social motivations. Overall, the market share of bounded rationality increased by 20 percentage points (a 100% increase).

These trends almost certainly underestimate the true magnitude of the cognitive turn. First, many of the papers that are discussed in this review are currently in the publication process. Second, many of the papers from the 2000s and 2010s that I code as bounded rationality do not share this review’s focus on cognitive foundations per se.

Objectives and main messages. Given the rapidly-growing body of research on the topic, the empirical results can arguably appear messy and far from suggesting the existence of common threads. One of my main objectives is to put some structure on the recent flurry of empirical results by categorizing them and relating them to the extant theory literature. Some of the topics covered in this review have received attention in more specialized and selective prior surveys, most of which focus on a particular strand of the theory literature (Maćkowiak et al., 2023; Gabaix, 2019; Bordalo et al., 2022c; Woodford, 2020; de Clippel and Rozen, 2023), with the exception of reviews on inattention (Loewenstein and Wojtowicz, 2023; Handel and Schwartzstein, 2018) and contingent reasoning (Niederle and Vespa,

2023). The present survey goes beyond these existing ones in two main respects. First, I focus on synthesizing the empirical and experimental evidence across the different strands of the cognitive foundations literature, and to draw out their commonality of imperfections in cognitive information processing. Second, none of the aforementioned reviews cover empirical work on complexity, cognitive noise and memory, which are among the most active literatures on cognitive foundations at this point.

The review will emphasize the following points:

1. Most economic decisions are difficult because they require intensive information processing. This includes the processes of attending to and remembering large amounts of information, but also the cognitive act of aggregating and trading off the relevant problem dimensions to reach a decision – even very low-dimensional problems are often difficult when they involve pronounced tradeoffs.
2. To reduce information-processing demands, people rely on broad classes of simplification strategies that play out in very similar ways across different contexts and domains. These include (i) noisy approximations and resulting behavioral attenuation; (ii) comparative thinking; (iii) reducing cardinality by overweighting what's salient, gets cued in memory, or is deemed important; (iv) thinking in analogies and categories; and (v) devaluing or shying away from difficulty-to-assess options.
3. We are beginning to understand how features of the decision problem (such as its complexity) affect the magnitude of information-processing imperfections.
4. Insights from the cognitive turn have started trickling down into applied fields, and have illuminated contexts in fields as diverse as finance, labor, and development. However, a broader integration of the cognitive turn into applied (behavioral) economics is only beginning.
5. Many classic behavioral anomalies reflect special cases of the five simplification strategies. The field has thus managed to partially consolidate a (long) laundry list of empirical regularities into a (shorter) laundry list of simplification strategies. However, we do not yet have a unified model that brings together the different simplification strategies, or even derives them from common principles.

Two caveats are in order. First, while this review is written under the basic presumption that economists stand to gain from studying information-processing imperfections, I am (very) far from asserting that people's preferences are as simple as the canonical economic model assumes. There can be little doubt that many empirical regularities in behavioral economics are not primarily driven by mistakes but reflect genuine preferences that are richer than the neoclassical model assumes. At the same time, this review will also emphasize the distinction between true and choice-revealed preferences, and that a great deal of behaviors

that we used to attribute to preferences actually reflect cognitive errors that are amenable to unification.

A second caveat is that this review is more of a progress report than a post-hoc synthesis. The cognitive turn is well underway but far from complete. Nonetheless, I believe we have arrived at a juncture where reviewing and tying together different lines of work is a productive endeavor to undertake, both to synthesize and to highlight open questions.

Organization of this review. It is common for behavioral economics reviews to be organized by decision domains (e.g., beliefs, choice under risk, intertemporal choice etc.). I deviate from this practice to highlight what I believe to be a key takeaway from the cognitive turn: that the same psychological principles underlie behavior across different decision domains. As I discuss in Section 2, the cognitive turn has emphasized the important role of information processing in forming mental representations and computing optimal decisions. Sections 3–8 are structured around the simplification strategies people use to ease information processing demands. The downside of this organization is that readers with an interest in a particular application cannot easily locate the relevant discussions because they will be scattered across the different simplification strategies. Thus, in Section 9, I revisit some of the main behavioral economics classics and explain whether and how the cognitive turn has changed our understanding of them. While most of the work in this review is experimental or theoretical in nature, Section 10 illustrates the potential promise of cognitive foundations for applied work. Finally, Section 11 discusses what I perceive to be open questions and methodological challenges.

2 Overview: Information Processing and Simplification

2.1 Policy Uncertainty

Consider a decision maker who is tasked with taking a decision a to maximize overall utility, which is comprised of different problem dimensions. Here, i indexes dimensions and γ_i scales the relative importance of each dimension. Dimension-by-dimension utility, $u_i(y_i)$, is a function of outcomes, y_i . Outcomes are produced as functions, $g_i(\cdot)$, of the decision and a vector of economic fundamentals, denoted θ , with elements θ_j :

$$\max_a U = \sum_i \gamma_i u_i(y_i) = \sum_i \gamma_i u_i(g_i(a, \theta)). \quad (1)$$

For example, in a labor supply context, the dimension-by-dimension outcomes y_i could be earnings and leisure that are implied by a function of hours worked (a) and the wage (θ_j), with γ_i scaling the relative importance of consumption and leisure. Similarly, in a savings context, y_i is consumption in period i , which is a function of savings, a , and the interest

rate, θ_j , with γ_i capturing intertemporal discounting.

The decision maker would like to choose the action that maximizes utility (whether or not utility includes non-standard elements such as social preferences):

$$\underset{a}{\operatorname{argmax}} U = a^* = f(\theta), \quad \left. \frac{\partial a^*}{\partial \theta_j} \right|_{\theta} \equiv \beta_j(\theta) \quad (2)$$

In the optimal policy function $f(\cdot)$, the vector of β_j represents the decision maker's *optimal decision weights*, by which I mean those weights that locally map economic fundamentals into the true utility-maximizing decision.

More generally, this optimal policy function need not refer to a choice problem. Rather, it could also apply to assessments of quality, or inference and forecasting problems. For instance, suppose the decision maker receives multiple signals about candidate quality, θ_j , that need to be aggregated into an overall assessment, or that the decision maker predicts the future realization of an autocorrelated process based on past observations, θ_j . In these cases, the normative weights, β_j , that optimally map observed fundamentals into a decision would correspond to Bayesian updating.

Whenever we use the rational actor model, we implicitly assume that people know how to translate any given set of primitives (their preferences, constraints and information) into the ex-ante optimal decision. In reality, people of course do not always seamlessly know how to map fundamentals into optimal decisions. Consider the following examples, some of which are ecological in nature and some of which represent stylized lab settings.

1. Taking as given your preferences and your beliefs about how effort at work affects your future promotions and salary, how many hours should you work per week to maximize your discounted expected lifetime utility?
2. Suppose the interest rate increases from 0% to 4%. What additional fraction of your income should you save given this change in circumstances?
3. Taking as given your return expectations and your risk preferences, which equity share maximizes your expected utility?
4. Taking as given your social motivations, how much should you donate to poor kids in developing nations to maximize your utility?
5. You are asked to evaluate PhD applicants on a scale from 0 to 10, where 0 means "probability of passing first-year courses is below 50%" and 10 means "probability of winning Clark medal is greater than 5%". An applicant has a GPA of 3.7 in Math and Econ at Chicago, a B+ in real analysis, a quantitative GRE score of 780, did a two-year pre-doc at Princeton (you have the letter, which says the candidate is an 8), and wrote a thesis that won a departmental prize. What's your rating?

6. Do you prefer a 30% chance of getting \$120 or a 85% chance of getting \$40?

Arguably, all of these examples trigger *policy uncertainty*: we do not know how exactly to translate the economic fundamentals into a utility-maximizing decision (or into an optimal policy function). This raises two immediate questions. First, why don't people know β_j ? Second, what do they do instead?

2.2 Information Processing

Economic decision problems routinely require intensive information processing, by which I mean attending to, remembering, aggregating and trading off economic variables to reach a decision.³ In casual terms, economists and cognitive scientists sometimes partition information processing into two components. First, mentally representing a problem – building “mental models” that ascertain which features and causal relationships exist. Second, computationally solving a problem – combining the different considerations that are part of one's mental model into a decision. Both of these steps require cognitive acts of information processing. Building a mental representation requires attending to and remembering those aspects that are relevant for the problem at hand. Similarly, computations require aggregating and trading off different problem dimensions – people cannot simply “pull out the max operator” like in a micro problem set. Instead, they need to implement some cognitive procedure to assess value, compute consequences and identify the most promising course of action. A main reason why economic decision problems often involve a cognitive act of aggregation is that carriers of utility (or of expected or discounted utility) are typically disaggregated – they consist of multiple components that jointly determine overall value. For example, determining the discounted utility of a savings plan requires aggregating multiple consumption events and one's discount function.

The literature has emphasized two broad classes of factors that affect the magnitude of information processing demands.

Cardinality. Adequately representing and solving a problem typically requires more intensive information processing when “more bits” are required to describe a problem – when more variables affect the outcome of interest, when the choice set has higher cardinality, when products have more dimensions, when more information signals need to be processed, and so on. Indeed, a broad body of evidence supports the idea that cardinality affects complexity and resulting information processing demands. For example, in lottery choice, people's behavior depends on the number of distinct payout states (e.g., Huck and Weizsäcker,

³This definition of information processing is different from the ones that are based on the neoclassical rational actor model. In that framework, “processing information” usually refers to situations in which people have uncertainty about some stochastic variable and update their beliefs upon receipt of information (for example about whether or not the stock market will go up next year). In contrast, I here reference a broader definition of information processing that even applies to entirely deterministic settings.

1999; Sonsino et al., 2002; Puri, 2022; Arrieta and Nielsen, 2023; Enke and Shubatt, 2023) or the menu size (Iyengar and Kamenica, 2010; Dertwinkel-Kalt and Köster, 2020b; Carvalho and Silverman, 2024). Similarly, people’s intertemporal decisions systematically vary as a function of how disaggregated future payments are across multiple dates (Dertwinkel-Kalt et al., 2022), and effort supply depends on the number of marginal tax rates (Abeler and Jäger, 2015). In the literature on belief updating and mental models, the presence of biases and heuristics increases in the number of information pieces that need to be aggregated (Enke and Zimmermann, 2019), the number of distinct states of the world (Ba et al., 2022), the number of causal nodes in a directed acyclic graph (Kendall and Oprea, 2022), the number of nodes in a game tree (Salant and Spenkuch, 2022) or the number of components of a strategic disclosure report (Jin et al., 2022). Building on some of these insights, Puri (2022) and Gabaix and Graeber (2023) propose models that formalize how the (importance-weighted) number of components of a problem drives complexity and shapes behavior.

Aggregations and tradeoffs. Aggregating the different components of a problem into a decision is usually more difficult when there are more components. Perhaps because of the obvious insight that processing costs and difficulty often increase in the size of a problem, my impression is that we as behavioral economists have developed a knee-jerk reaction to interpret non-standard behavior in “simple” low-dimensional problems as reflecting non-standard preferences or systematic biases. In fact, much experimental work is explicitly motivated by the desire to document some effect of interest in a setting that is as simple as possible, supposedly ruling out information processing costs as a source of “non-standard” behavior.

Yet, building on a rich literature in psychology (e.g., Tversky and Shafir, 1992; Drugowitsch et al., 2016), a main insight from much recent work is that even low-dimensional problems require intensive information processing, in particular when people need to trade off the relative advantages and disadvantages of different problem dimensions. For example, labor supply involves trading off money and leisure; belief updating requires trading off a prior and a signal that may point in different directions; consumer choice requires trading off prices and multiple dimensions of quality; choice under risk calls for trading off risk and expected return; intertemporal choice requires trading off delays and rewards; strategic decisions in extensive-form games require trading off the costs and benefits of different nodes of the game tree; and so on.

These cognitive acts require intensive information processing, for two reasons. First, information processing may be required for introspection: people may not know their true discount factor, their altruism parameter, their utility weight on leisure versus consumption, and so on. Put differently, people often don’t really know how to put different problem dimensions into a “common currency” (see Walasek and Brown, 2021, for a recent overview

in psychology). Second, even if people perfectly knew their preferences, information processing is required to aggregate the different problem dimensions into an optimal policy function. For instance, even conditional on knowing that one’s intertemporal discount factor is $\delta = 0.97$, it is entirely non-trivial to determine for how many hours to go to the gym this week.⁴

There is now a broad body of evidence that people find many low-dimensional economic decision problems difficult, and that the difficulty of navigating tradeoffs is a main reason for this. Decision theorists and experimentalists have provided converging evidence that people often struggle in ranking alternatives, a phenomenon referred to as (revealed) incomplete preferences (e.g., Ok et al., 2012; Halevy et al., 2023). Much of this evidence is derived from people explicitly acknowledging that they don’t know which option to choose. This class of experiments includes unincentivized self-reports that indicate choice uncertainty (Cohen et al., 1987; Loomes and Sugden, 1995; Cubitt et al., 2015; Bayrak and Hey, 2020; Enke and Graeber, 2023; Enke et al., 2023a, 2024a), documentations that people delegate the decision to a computer that estimates their preferences (Nielsen and Rigotti, 2022), that people reverse their decisions when confronted with axioms they endorsed (Nielsen and Rehbeck, 2022), or that they deliberately randomize (Agranov and Ortoleva, 2017, 2020; Agranov et al., 2020).⁵ In situations like these, standard revealed preferences techniques are problematic. For instance, building on the coherent arbitrariness work (Ariely et al., 2003), deClippel et al. (2024) document that choices can appear highly internally coherent and pass commonly-used tests for the existence of a stable underlying utility function, even when the decisions are very far from actually maximizing the decision-maker’s objective.

Revealed incomplete preferences almost certainly reflect the difficulty of navigating tradeoffs. After all, few people report difficulty with one-dimensional choice problems such as choosing between \$8 and \$10 – instead, people only struggle when they need to aggregate across dimensions that have different advantages and disadvantages. Indeed, multiple studies have documented that experimental manipulations of the strength of across-dimension tradeoffs have large effects on self-reported uncertainty and randomization (Tversky and Shafir, 1992; Agranov and Ortoleva, 2017; Enke and Shubatt, 2023; Shubatt and Yang, 2024).

The terminology of “incomplete preferences” is perhaps slightly misleading as a general primitive underlying policy uncertainty because it may be interpreted as saying that preference uncertainty is the sole driver of uncertainty. This is not the case. Rather, people also struggle with entirely objective low-dimensional problems that are specially designed to understand the complexity of tradeoffs and aggregation (Martínez-Marquina et al., 2019;

⁴A growing body of work emphasizes that problems tend to be more complex when the mapping between economic primitives to optimal decisions is non-linear (Rees-Jones and Taubinsky, 2020; Agranov and Reshidi, 2023).

⁵Deliberate randomization is also consistent with convex preferences, yet it strongly varies with problem difficulty (Agranov and Ortoleva, 2017), suggesting it is a response to information processing costs.

Oprea, 2022; Vieider, 2022; Enke et al., 2023a; Enke and Shubatt, 2023; deClippel et al., 2024). This suggests that a non-trivial part of imperfect information processing reflects the cognitive difficulty of aggregating known primitives into a decision.

Implications of an information processing account. Relative to accounts that emphasize non-standard preferences or fixed parametric biases, an information processing view has two immediate implications. First, the presence and magnitude of “non-standard” choice and belief patterns should depend on aspects that affect the quality and focus of information processing, including: (i) the problem’s complexity because it determines how much information needs to be processed; (ii) experience and the availability of cognitive resources because they determine how much information can actually be processed; and (iii) salience, memory cues and incentives because they determine which information is processed (first).

A second basic implication of an information processing perspective is that behavioral anomalies (even those that we typically attribute to non-standard preferences) should coincide with independent proxies for imperfect information processing, such as measures of decision noise, self-reported uncertainty, response times, eye movements, and others.

2.3 Simplification Strategies

What do people do when the information processing required to produce the correct problem representation and / or to trade off different problem components, is difficult or costly? Table 1 summarizes the main simplification strategies that I will discuss in this review. As I point out below, some of these literatures are more developed than others.

The first two simplification strategies are mostly targeted at simplifying the process of aggregating tradeoffs across different dimensions. First, instead of precisely aggregating, people noisily approximate. This produces *behavioral attenuation*: observed decisions and beliefs are often insufficiently elastic to variation in economic fundamentals. Second, instead of aggregating tradeoffs across dimensions, people bracket and engage in comparative thinking: they compare outcomes within each problem dimension. Both of these simplification strategies assert that people process all available variables, but do so in a noisy or comparative way.

A third class of simplification strategies, on the other hand, presumes that people work with a subset of variables: they ignore some and overweight others. This happens through a combination of low-level attentional processes (focusing on what stands out or is surprising), associations in memory (relying primarily on those experiences that are similar to the current experience), and goal-oriented processes (focusing on what’s deemed important).

Fourth, when a problem is sufficiently difficult, people may not even work with (a subset of) the correct variables but instead solve a different, related problem that they know how to solve. This can be thought of as “coarsening” reality and lumping together similar situations,

Table 1: Categories of simplification strategies

Simplification strategy	Empirical regularities	Classes of models
1. Noisily approximate	Behavioral attenuation; spurious anomalies	Noisy policy functions; Noisy perceptions
2. Bracket and compare	Reference points; contrast effects	Reference-dependence; normalizations
3. Reduce cardinality / work with subset	Overweight what’s deemed important, stands out or gets cued	Goal- or cue-directed attention and memory
4. Solve a related problem	Categorize / lump similar situations together	Analogies; associations
5. De-value or shy away	Complexity aversion	Complexity aversion; caution

decisions, or variables into categories or analogy classes.

Finally, people may simply shy away from options that require too much information processing to be properly evaluated.

High- and low-level simplification strategies. The terminology “simplification strategies” is not intended to imply that these are all necessarily deliberate and / or conscious in nature. While some likely are (such as noisily approximating), others likely aren’t (such as overweighting what stands out). I still view these as simplification strategies because – as discussed in greater detail below – it is conceivable that we have evolved tendencies to reduce information processing costs by following certain low-level cues. In this sense, these behaviors serve to simplify a problem, deliberate or not.

Nonetheless, the distinction between somewhat-deliberate simplification strategies and those that reflect low-level cues merits attention. Each of these two lines of work builds on a classic research tradition. Research on deliberate simplification strategies is often inspired by the bounded rationality movement pioneered by Simon (1956), which emphasizes complexity and bounds on computational capabilities. Research on misleading intuitions and gut feelings, on the other hand, is often inspired by the distinction between “system 1” and “system 2” thinking (Kahneman, 2011) – system 1 is said to be a fast-moving decision process that reflects gut instincts and responds to contextual cues, while system 2 is a slower and more deliberative system that potentially overrides misleading gut reactions.

Both of these approaches are needed because, as far as their applicability to economic decisions is concerned, each tradition in isolation arguably suffers from different weaknesses. The Simon movement with its emphasis on the costliness of cognitive procedures tends to

ignore that mistakes are often driven by contextual cues and associations. The “system 1” movement, on the other hand, acquired fame by constructing clever – but ecologically largely counterfactual – thought experiments that tended to suggest that the only real problem with human cognition was to get the intuitive system 1 under control. Yet as far as economic applications are concerned, this perspective is misleading because most economic decisions – looking for a job, investing in the stock market, buying a house – are very difficult even conditional on system 2 overriding system 1.

3 Noisy Approximations

A first main implication of imperfect information processing is noise – people approximate rather than solve a problem precisely. Perhaps because of the idea that noise often “cancels out”, it is sometimes treated almost as a nuisance in behavioral economics, as something that one needs to tack on to a behavioral model to actually make it work in empirical applications but that is not in itself worth studying. However, a key insight of much recent research is that *noise causes bias*. How can unsystematic noise cause bias? The literature has emphasized two such mechanisms. First, as discussed in Section 3.3, even mean-zero noise can cause bias when researchers look at “asymmetric” problem setups in which noise can only push in one direction. Second, even when the problem setup is entirely symmetric, noise can cause bias by producing *diminishing sensitivity*.

Diminishing sensitivity pervades many different literatures in behavioral economics. It refers to the observation that people’s choices, valuations, beliefs and expectations generally become more and more insensitive (“attenuated”) with respect to problem parameters (fundamentals) as the parameter moves further away from the “boundary” of the parameter space, in particular when these boundary points render a decision trivial (e.g., a price of zero or a payout probability of one).

In recent years, researchers have made much progress (i) in documenting how pervasive such “behavioral attenuation” and diminishing sensitivity are across very different contexts and (ii) in showing that these patterns reflect a generic implication of imperfect and costly information processing, rather than domain-specific preferences. This line of research builds on much earlier work on diminishing sensitivity, such as in the context of prospect theory or hyperbolic discounting (e.g., Kahneman and Tversky, 1979; Loewenstein and Prelec, 1992; Prelec and Loewenstein, 1991). The main difference is that the cognitive turn highlights information processing demands as underlying unifying principle, rather than attributing attenuation to unrelated domain-specific heuristics or preferences.

3.1 Behavioral Attenuation

It is instructive to initially focus on behavioral attenuation (or insensitivity): the observation that decisions are insufficiently sensitive to variation in parameters, relative to what neoclassical theory predicts. Through the lens of equations (1) and (2), attenuation means that the effective decision weight $\hat{\beta}_i$ satisfies $\hat{\beta}_i < \beta_i$. Attenuation is different from underreaction. Underreaction is defined for a single decision (it concerns the *level* of decisions), while attenuation is defined over at least two decisions (the *slope* of decisions). For example, in belief updating, attenuation to the precision of information mechanically produces overreaction to weak signals but underreaction to strong signals (Augenblick et al., 2021). To take a different example, when people forecast an AR(1) process and are attenuated to the degree of autocorrelation, then expectations will mechanically overreact to the last observation whenever the true autocorrelation is relatively small (Enke et al., 2024a).

According to an information processing view, behavioral attenuation is akin to attenuation bias in econometrics. The main difference is that it does not arise because the fundamental θ is measured with noise but, instead, because the difficulty of translating θ into an optimal decision produces noise in the decision maker’s mind.

3.1.1 Evidence

Table 2 presents an overview of phenomena in the literature that can be understood as attenuation effects. In almost all cases, the attenuation is such that decisions are “too high” for some parameters and “too low” for other parameters. This canonical “flipping” property helps to distinguish behavioral attenuation from directional biases. Consider the following illustrative examples:

- According to the the probability weighting function, people are more risk seeking than the expected utility model prescribes when the payout probability of a gain is small, but less risk seeking than the benchmark when the payout probability is large – an insensitivity to the probability.
- According to the hyperbolic discounting function over financial flows, people act as if they are less patient than the exponential model prescribes when the delay is short, but more patient when the delay is very long – an insensitivity to the delay.
- In belief updating, people update too much when the information is very imprecise but too little when it is very precise – an insensitivity to signal precision.
- People’s prediction of their income tax burden is too high when their income is low but too low when their income is high – an insensitivity to the tax rate.
- In medical testing, doctors overttest when patient risk is low but undertest when patient risk is high – an insensitivity to patient risk.

- People’s stock market forecasts are too optimistic when they are polled about objectively unlikely events but too pessimistic when they are polled about objectively likely events – an insensitivity to probabilities.

These examples are just the tip of the iceberg – attenuation is found in lab experiments that involve people’s preferences (choice contexts), in laboratory belief updating or prediction problems that have objectively correct solutions, in strategic games, in field experiments, in large-scale expectations surveys and in observational data.

Evidence from the lab. Enke et al. (2024a) implement a large set of experiments that were crowd-sourced from independent experts. These experiments range from preference elicitation to belief updating to generic optimization problems, involving risk, time, consumption-savings, effort supply, taxes, fairness, prediction and inference, in both individual decisions and strategic games. In each experiment, they measure cognitive uncertainty and link it to the elasticity of decisions as a function of problem fundamentals. They find significant behavioral attenuation in almost all tasks, meaning that in nearly all tasks those decisions that are associated with higher cognitive uncertainty are more attenuated functions of the relevant economic primitive. Because cognitive uncertainty is a proxy for the magnitude of information-processing imperfections, this is interpreted as evidence for behavioral attenuation.

Perhaps the most common application of the idea of information processing- and noise-driven attenuation so far is in the domain of choice under risk. First, recent work has highlighted that the insensitivity of prospect theory’s probability weighting function is not driven by preferences but rather by the difficulty of aggregating payouts and probabilities (Enke and Graeber, 2023; Oprea, 2022; Frydman and Jin, 2023; Vieider, 2022).

Second, the insensitivity of decisions to objective risk that is captured by the probability weighting function has a direct counterpart in the literature on choice under ambiguity. Recent studies and literature reviews highlight the notion of so-called “a-insensitivity” (for ambiguity insensitivity), see Dimmock et al. (2015), Trautmann and Van De Kuilen (2015) and Henkel (2022). According to this literature, people’s decisions are excessively insensitive to whether the likelihood of an ambiguous event is described as relatively high or relatively low – thus a version of the probability weighting function in the presence of Knightian uncertainty. This is a form of attenuation, albeit with respect to ambiguous likelihoods rather than objective probabilities.

A third application in the literature is to the “attenuation puzzle” in household finance, which refers to the strongly attenuated link between people’s investment decisions and their subjective return expectations. This attenuation, too, appears to reflect the difficulty of aggregating problem components into a decision (Yang, 2023; Charles et al., 2022).

The link between information processing and attenuation is not restricted to choice under risk or uncertainty. Consider the canonical hyperbolic discounting function over money

Table 2: Empirical attenuation patterns across domains

Decision	Attenuation wrt	Label in literature	Examples
<u>Choice contexts in the lab</u>			
Choice under risk	Payout probability	Probability weighting	Enke and Graeber (2023); Oprea (2022)
Choice under risk	Payment	Non-linear value function	Frydman and Jin (2022); Khaw et al. (2021)
Choice under ambiguity	Likel. of ambig. event	Ambiguity insensitivity	Trautmann and Van De Kuilen (2015)
Portfolio allocation	Subj. expected return	Attenuation puzzle	Charles et al. (2022); Yang (2023)
Time discounting	Time delay	Hyperbolic discounting over money	Enke et al. (2023a); Ebert and Prelec (2007)
Extended dictator game	Group size	Insensitivity to group size	Schumacher et al. (2017)
Attention allocation	Stake size	Increasing Shannon cost	Dean and Neligh (2023); Bronchetti et al. (2023)
Product demand	Tax rate	Insensitivity	Morrison and Taubinsky (2019)
Labor supply	Tax rate	Attenuation to tax rate	Abeler and Jäger (2015)
Forward reasoning	Contingent payouts	Cognitive attenuation	Chakraborty and Kendall (2022)
Communication	Verifiability of info	Over- and under-communic.	Fréchette et al. (2022)
News vendor game	Marginal cost	Pull-to-center effect	Schweitzer and Cachon (2000)
<u>Beliefs in the lab</u>			
Belief updating	Signal strength	Under- and overreaction	Augenblick et al. (2021)
Belief updating	Base rate	Base rate insensitivity	Enke and Graeber (2023); Benjamin (2019)
Belief updating	Uncertain information	Uncertainty-induced insensit.	Liang (2023b)
Information demand	Signal strength	Compression effect	Ambuehl and Li (2018); Liang (2023a)
Confidence	True ability	Hard-easy effect	Moore and Healy (2008)
Gender stereotypes	True ability	Difficulty-influenced miscal.	Bordalo et al. (2019)
Forecasting	Persistence of process	Overextrapolation	Afrouzi et al. (2023)
Mental model	Narratives / DAGs	Compression to middle	Kendall and Charles (2022)
<u>Strategic games in the lab</u>			
Investment coordination	Value of project	Insensitivity to payoffs	Frydman and Nunnari (2023)
Order statistics in auctions	Signal rank		Nagel et al. (2023)
<u>Field experiments</u>			
Contingent valuation	Effectiveness of policy	Scope insensitivity	Toma and Bell (2022)
Perception of tax rate	Tax rate	Scheduling	Rees-Jones and Taubinsky (2020)
Labor supply	Ratchet incentives	Shrouding	Abeler et al. (2023)
<u>Large-scale surveys on beliefs and perceptions</u>			
Subjective income rank	True income rank	Pull-to-center	Hvidberg et al. (2023)
Stock market expectations	Historical probabilities	Pull-to-center	Enke and Graeber (2023)
Stock market expectations	Time interval	Sub-additive time perception	Haan et al. (2022)
Valuation of health states	Health	Attenuation	Kaats (2023)
2nd-order policy beliefs	True support for policy	Minority salience	Bursztyn et al. (2023)
<u>Observational field data</u>			
Equity share	Subjective returns	Attenuation puzzle	Drerup et al. (2017); Giglio et al. (2021)
Medical testing	Patient risk	Over- and undertesting	Mullainathan and Obermeyer (2022)
Borrowing	Policy rates	Unresponsive to incentives	D'Acunto et al. (2023)
Price setting	Local demand	Uniform pricing	DellaVigna and Gentzkow (2019)
Wage setting	Local supply	National wage setting	Hazell et al. (2022)
<u>Evidence in psychology (perceptual decision tasks)</u>			
Magnitude estimation (many)	True magnitude	Central tendency effect	Petzschner et al. (2015); Xiang et al. (2021)

Notes. Papers that document behavioral elasticities that are commonly viewed as insufficiently large. Whenever possible, the list of references prioritizes those papers that directly link the relevant phenomenon to complexity or noise. Additional references are in the main text.

and the widely-known “scope insensitivity” effect in contingent valuation studies. Both of these effects boil down to the pattern described above: away from the boundary of zero, decisions are very insensitive to variation in the delay (in intertemporal discounting) or the scope of a policy proposal such as the number of birds saved (in contingent valuation studies). Recent work has shown that both of these attenuation patterns are driven by imperfect information processing (Ebert and Prelec, 2007; Enke et al., 2023a, 2024a; Toma and Bell, 2022).

An important indication that behavioral attenuation primarily reflects imperfect information processing is that it is pervasive both in choice contexts that involve people’s preferences and in situations in which an objectively correct decision exists, such as belief updating, forecasting, and information acquisition problems. For instance, people’s beliefs are heavily attenuated functions of the precision of the information they received. This “likelihood insensitivity” means that if people receive imprecise signals, they update too much and if they receive precise ones, they update too little (Enke and Graeber, 2023; Augenblick et al., 2021; Ba et al., 2022; Prat-Carrabin and Woodford, 2022; Ambuehl and Li, 2018).

To support the idea that behavioral attenuation reflects imperfect information processing, researchers have made use of a number of different methodological tools. First, attenuation effects are typically strongly correlated with cognitive uncertainty: people’s self-reports of how certain they are that their decision reflects what they actually prefer (Enke and Graeber, 2023; Enke et al., 2023a, 2024a; Giglio et al., 2021; Yang, 2023; Augenblick et al., 2021). Second, exogenous complexity manipulations consistently produce more pronounced attenuation (Enke and Graeber, 2023; Enke et al., 2023a; Gabaix and Graeber, 2023). Third, attenuation persists when the choice problems are converted into objective problems, removing scope for non-standard preferences (Oprea, 2022; Vieider, 2022; Enke and Shubatt, 2023). Fourth, attenuation is locally reduced as people garner more experience with the problem (Frydman and Jin, 2022, 2023; Frydman and Nunnari, 2023). Fifth, attention is more pronounced when people pay less attention to the relevant variable (e.g., Ebert and Prelec, 2007; Hartzmark et al., 2021; Imas et al., 2022) and when their cognitive resources are lower or exogenously depleted (e.g., Ebert and Prelec, 2007; Choi et al., 2022; Enke and Graeber, 2023).

In summary, there is now a large body of evidence from the lab that supports two broad propositions. First, behavioral attenuation is very widespread. Second, attenuation is empirically tightly linked to imperfect information processing.

Evidence from surveys of expectations and perceptions. A large and growing literature measures people’s expectations and perceptions about economic quantities. A first-order stylized fact that emerges from this literature is that there are very large attenuation effects: subjective assessments are a strongly compressed function of objective quantities. For instance, Hvidberg et al. (2023) show that people’s perception of their income rank

is a strongly compressed function of their actual rank (“everyone thinks they are middle class”). This compression pattern is strongly correlated with self-reported cognitive uncertainty (Enke and Graeber, 2023). Similarly, people’s stock market return expectations or their inflation expectations are heavily attenuated functions of true (historical) probabilities, which is again correlated with cognitive uncertainty.

Evidence from the field. There is an emerging body of evidence that points to the importance of behavioral attenuation also for field decisions, though there is less evidence to date that directly ties it to imperfect information processing. For example, the link between subjective return expectations and people’s equity shares is heavily attenuated relative to theoretical benchmarks such as the Merton model (Drerup et al., 2017; Ameriks et al., 2020; Giglio et al., 2021). Relatedly, the elasticity of spending to inflation expectations (D’Acunto et al., 2019) or of medical testing to patient risk (Mullainathan and Obermeyer, 2022) are substantially attenuated. As a result, doctors overtest when patient risk is low but undertest when patient risk is high – the canonical “flipping” pattern that is associated with behavioral attenuation.

Recent evidence suggests that these attenuation patterns also affect responses to policy: consumers’ response to the incentives resulting from changes in government programs is heavily attenuated, which affects behaviors such as borrowing and car purchases (D’Acunto et al., 2023). Likewise, people’s perception of their tax schedule shows strong attenuation effects: as the true marginal tax rate increases, people’s perceived marginal rate also increases, but at an attenuated rate, meaning that people overestimate the tax burden at low incomes and underestimate it at high incomes (Rees-Jones and Taubinsky, 2020; Abeler and Jäger, 2015).

There is now also evidence that workers’ labor supply is a heavily attenuated function of the incentives that are implicit in the contract. Abeler et al. (2023) present a combination of field and online experiments with warehouse workers in which they document that workers’ response to the presence of dynamic Ratchet incentives is heavily attenuated due to the complexity of the incentive contract.

Evidence in cognitive psychology. Psychologists have long been aware that attenuation is “everywhere” also in perceptual decision making. For example, in estimating angles, distances, elapsed time or the number of dots on a screen, people typically overestimate relatively small quantities but underestimate relatively large ones. When Kahneman and Tversky (1979) initially motivated prospect theory, they did so with explicit reference to this work on perception. See Petzschner et al. (2015) for a recent review of behavioral attenuation in magnitude estimation tasks, and how these can be understood through Bayesian models of cognitive noise.

3.1.2 Models of Behavioral Attenuation

Multiple models have been proposed to capture behavioral attenuation, including its link to measures of noise. In purely interpretive terms, the idea behind various models is that the cognitive act of trading off and aggregating multiple problem dimensions produces noise, and that this noise generates attenuation.

Loosely speaking, many models generate attenuation through a reduced-form equation that resembles the classical anchoring-and-adjustment heuristic. To illustrate, suppose that people’s average decisions can be described by a convex combination of a fixed default (d) and the true utility-maximizing decision $a^*(\theta)$ at the current parameter vector θ :

$$E[a^o(\theta)] = (1 - \lambda) a^*(\theta) + \lambda d, \quad (3)$$

Here, d could reflect a kind of prior, while $(1 - \lambda) \in [0, 1]$ captures the decision-maker’s level of rationality that governs the degree of attenuation. This simple reduced-form decision rule performs surprisingly well in proximally explaining many of the attenuation patterns enumerated above. To see the attenuation effect, notice that

$$\frac{\partial E[a^o(\theta)]}{\partial \theta} = \underbrace{(1 - \lambda)}_{\text{Attenuation } (<1)} \underbrace{\frac{\partial a^*(\theta)}{\partial \theta}}_{\text{Normative sensitivity}}, \quad (4)$$

where by “normative sensitivity” I mean the sensitivity the decision maker would exhibit if they were to maximize their true objective function (whether or not that objective function includes non-standard preferences).

Various modeling approaches take a stance on how specific cognitive foundations generate variants of this decision rule. While these models differ in the details, their underlying logic is often very similar. We can distinguish between two questions here. First, why does attenuation emerge and what determines its magnitude (λ)? Second, what determines the default (d)?

On the first question, I discuss two classes of models that have recently attracted attention: (i) models of *policy uncertainty*, in which the decision maker does not know the function $a^*(\theta)$ that maps observed problem parameters into the utility-maximizing decision; and (ii) models of *noisy perceptions* (or noisy parameters), in which the decision maker exhibits a noisy perception of the parameter θ .

Both of these classes of models are inspired by Bayesian models of noisy perception in decision neuroscience. These models are typically applied to situations in which a decision maker must estimate a magnitude – an angle, the number of dots on a screen, and so on. The decision maker holds a prior about the quantity (say, the average number of dots on the screen across trials in the experiment), and generates a noisy perceptual signal by inspecting the quantity (say, by briefly looking at the dot cloud). The decision maker is then assumed

to combine his prior and the perceptual signal in a Bayesian way, which often gives rise to attenuated decision rules of the form in eq. (3).

When applied to economic decisions, these models are typically either tweaked or re-interpreted. As discussed above, the main difficulty in economic choice is usually not the literal perception of a quantity. Rather, cognitive noise emerges because mapping (or aggregating) a given set of problem fundamentals into a decision is difficult.

Models of policy uncertainty. Ilut and Valchev (2023) propose a model that formalizes the idea that people do not know the function $a^*(\theta)$ that maps the observed fundamental θ into the utility-maximizing decision. The decision maker perfectly observes the fundamental (or state) θ but does not know how to map it into a decision. Through deliberation, the decision maker generates noisy cognitive signals about the optimal decision at the prevailing parameter value: $s(\theta) = a^*(\theta) + \epsilon$, and uses this cognitive signal to update about the policy function. This type of model – that explicitly features uncertainty about the optimal policy function – is an attractive way to think about why people express cognitive uncertainty and exhibit noisy decisions. A related model is presented by Yang (2023), in which decision makers exhibit uncertainty over the function that maps their subjective stock return expectations into their expected-utility-maximizing equity share. Also see the model in Enke and Graeber (2023).

Noisy perception models. The models summarized above are ones of *noisy policy functions*. Another technique to conceptualize the link between noise and behavioral attenuation is models of *noisy perceptions*, in particular Bayesian models of cognitive noise (Gabaix, 2019; Woodford, 2020). The key feature of these models is the assumption that the decision maker only has a noisy perception of (or pays only partial attention to) the – in principle known – parameter θ . For example, the decision maker may observe a noisy cognitive signal, $s = \theta + \epsilon$, and thus shrinks the true θ to some prior, akin to equation (3). These models generally generate attenuation effects. For instance, Khaw et al. (2021, 2022) and Vieider (2022) show how the probability weighting function can be derived by assuming a noisy perception of the payout probability, Woodford (2012b,a) and Frydman and Jin (2022) show that the attenuation of prospect theory’s value function follows from noisy perceptions of the payout amount, Lian (2021) derives narrow bracketing as an implication of noisy perception of prices, and Gershman and Bhui (2019), Gabaix and Laibson (2022) and Vieider (2021) all show how different versions of hyperbolic discounting follow from the idea that either future utils or the delay in an intertemporal decision problem are perceived with noise. Da Silveira et al. (2020) show that attenuation to the persistence parameter in forecasting can be captured through a noisy memory model.

A challenge in interpreting some of these “noisy parameter” or “noisy perception” models is that they are routinely applied to contexts in which the economic fundamental θ is ob-

jectively known and saliently displayed. Initially, researchers often interpreted “parameter noise” as capturing noisy low-level cognition such as misperception of numbers. However, as evidence both inside and outside of economics accumulated (e.g., Drugowitsch et al., 2016), it has become clearer that the vast majority of imperfect information processing and noise do not reflect the literal noisy perception of a number but rather the difficulty of translating that number into an optimal decision. In response, researchers have taken a probably more realistic and broader interpretation of “noisy parameter” models as formalizing the difficulty of aggregating different aspects of the problem into a decision.

Related to this discussion are also drift-diffusion models (Ratcliff, 1978; Ratcliff and McKoon, 2008; Krajbich et al., 2012; Fudenberg et al., 2018). Unlike the static Bayesian cognitive noise models, these are sequential sampling models in which the level of uncertainty at the moment of the decision partly reflects the length of deliberation. One way to interpret the static cognitive noise models is that the level of cognitive noise reflects such an underlying dynamic process, in which the decision maker sequentially accumulates information about the uncertain object and combines it into a posterior level of uncertainty.

Regardless of their specifics, the models summarized above generally generate the prediction – routinely confirmed in experimental data – that *noise and bias are linked*: the decision maker’s degree of attenuation is correlated with their noisiness. In the models reviewed above, bias emerges because of the shrinkage to the invariant default (or prior) d . This raises the question what this prior is and how it should be specified.

The cognitive default. One interpretation is that the prior corresponds to what “usually happens” – when we encounter a decision situation, we often have a sense for what we will do even before we actively deliberate about the problem. In this regard, a useful distinction is between contexts with which the decision maker has experience, and entirely unfamiliar environments (as is often the case in abstract lab experiments). In contexts with which people have experience, a plausible idea is that the cognitive default decision (the “prior”) may be shaped by memory: it is given by what one usually does (or did in the past). For example, when we commute to work, we don’t actively reason about the best route but simply follow a default strategy, unless the circumstances change in significant ways. There is less consensus in the literature about how the prior should be specified in situations with which the decision maker has no experience. Some evidence suggests that the default decision consists of appealing simple heuristics that “make sense on average”. For example, in two-state belief formation problems, a default decision of 50-50 is correct “on average” if the distribution of problems is symmetric (Enke and Graeber, 2023; Ba et al., 2022). Similarly, in preferential choice tasks, a default decision that reflects a “compromise effect” (for example, switching in the middle of an experimental price list) makes sense if one has average preferences in the population (e.g., Kamenica, 2008; Beauchamp et al., 2019).

However, overall, there is much we don’t know about what the right way of thinking

about cognitive defaults is.

Models in psychology. Just like in economics, psychologists explain behavioral attenuation patterns through a variety of noise models that first emerged in the 1990s, such as Bayesian models of noisy cognition (Chater et al., 2008; Griffiths et al., 2008; Gershman, 2021; Bhatia and Loomes, 2017; He et al., 2019), drift-diffusion models and models of reinforcement learning. This line of work is also closely related to models of resource-rational cognition that enjoy great popularity in cognitive science today (Lieder and Griffiths, 2020), of which rational inattention can be viewed as a special case. Interestingly, some of the main challenges that behavioral economists have encountered in working with these models appear in very similar ways in cognitive psychology. For example, (the lack of) the specification of the prior in Bayesian cognitive noise models – often criticized by economists – has produced heated debates also in cognitive psychology (e.g., Bowers and Davis, 2012).

Diminishing sensitivity. If imperfect information processing generates behavioral attenuation (insensitivity), then why is it that researchers commonly identify *diminishing* sensitivity, including occasional excess sensitivity at certain points? In short, the answer is that, for a variety of potential reasons, some decision problems produce more noise (and hence attenuation) than others, as I discuss now.

3.2 What Determines the Magnitude of Noise?

The strength of tradeoffs. A recurring theme in this review is that much of the difficulty of low-dimensional decision problems is driven by the need to aggregate tradeoffs across different problem dimensions (Tversky and Shafir, 1992). If true, this suggests that noisy approximations and behavioral attenuation should be more pronounced if aggregating tradeoffs is more difficult. Indeed, much evidence shows that people’s decisions are considerably less noisy in the neighborhood of dominance (e.g., Agranov and Ortoleva, 2017; Enke and Shubatt, 2023; Shubatt and Yang, 2024; Enke et al., 2024a). For example, trading off the costs and benefits of an additional hour of work is arguably cognitively difficult when the wage is, say, \$35 but it is entirely trivial when the wage is \$0. From this perspective, diminishing sensitivity is not a counterexample to the prevalence of behavioral attenuation – instead, it is a *result* of behavioral attenuation, or more specifically, a result of the fact that there is less behavioral attenuation in some problems than in others.

A growing body of theoretical and experimental work studies the more general idea that *the strength of across-dimension tradeoffs causes complexity* (and hence noise) because it makes comparisons difficult. Intuitively, holding fixed the aggregated value difference between two options, they are easier to compare when they are more similar to each other dimension-by-dimension because the decision maker does not have to deal with pronounced

tradeoffs across dimensions.

The idea that navigating across-dimension tradeoffs make choice difficult dates back at least to Tversky (1969). Rubinstein (1988, 2003) suggested that people cancel dimensions when the attributes are “similar”. More formally, Natenzon (2019) and He and Natenzon (2022, 2023) formalize the idea that dimension-by-dimension dissimilarity generates comparison complexity in multiattribute choice and show that it explains decoy effects. Subsequently, Shubatt and Yang (2024) develop a metric of comparison complexity that is applicable to multi-attribute, lottery and intertemporal choice, which intuitively formalizes the strength of across-dimension tradeoffs as “distance to dominance.” They show that this metric of complexity explains a wide range of evidence on the context-dependent nature of behavioral economics phenomena such as probability weighting, hyperbolic discounting over money, and systematic preference reversals, including why these anomalies are often found to depend on the elicitation method (Harbaugh et al., 2010; Andreoni and Sprenger, 2011; Bouchouicha et al., 2023).

A growing body of empirical work documents that tradeoff complexity (dimension-by-dimension dissimilarity) indeed strongly drives choice mistakes. Enke and Shubatt (2023) quantify the complexity of lottery choice problems and find that that by far the most important determinant of lottery choice complexity is dissimilarity: how disaggregated the overall expected values difference between two lotteries is across multiple states. Shubatt and Yang (2024) find similar patterns in multi-attribute and intertemporal choice. Psychologists have identified very similar patterns (Erev et al., 2010). The estimated effects of dissimilarity on choice errors are typically very large, suggesting that it is a first-order determinant of the noisiness of people’s decision process.

Optimizing considerations. A second category of approaches is that the degree of noisiness reflects optimizing considerations: the decision maker chooses how long to deliberate to trade off the benefits of a more precise cognitive representation against thinking costs. This broad category of models includes theories of rational inattention (Caplin and Dean, 2015; Caplin et al., 2020) and sparsity (Gabaix, 2014, 2023; Gabaix and Graeber, 2023). Both of these models predict that cognitive effort devoted to a problem (or a variable) increases in the problem’s importance, such that more important problems give rise to less noise. I do not discuss the details here because this literature has been reviewed elsewhere (Maćkowiak et al., 2023; Gabaix, 2019).

Part of this optimizing class of theories are also models of efficient coding (e.g. Friedman, 1989; Robson, 2001; Netzer, 2009; Woodford, 2012a,b; Netzer et al., 2022; Herold and Netzer, 2023; Steiner and Stewart, 2016; Frydman and Jin, 2022, 2023; Frydman and Nunnari, 2023; Khaw et al., 2022). The main idea in these models is that the decision maker chooses to have a more precise (less noisy) mental representation for problems that he encounters more often. These models very naturally generate diminishing sensitivity with respect to

frequently-encountered “reference points”. For instance, consider prospect theory’s value function, which exhibits diminishing sensitivity for both gains and losses. In most applications, the reference point is a point with which the decision maker has much experience. Then, efficient coding predicts that the value function (and decisions) are very sensitive to changes in circumstances around the reference point but insensitive far away from it.

Recent experimental work has applied this idea to lottery choice and strategic decisions (Frydman and Jin, 2022, 2023; Frydman and Nunnari, 2023). The main insight of this class of papers is that as people garner more experience with a certain neighborhood of problems, their decisions become more sensitive to variation in fundamentals within this neighborhood (but not outside of it). For instance, in choosing between risky assets, decisions become more sensitive to any given change in payoffs if the decision maker has encountered these payoffs more frequently in the past.

3.3 Asymmetric Problem Setups and Spurious Behavioral Motivations

A second reason for why recent research often documents that noise can cause bias is that prior work often provided evidence for non-standard preferences using discrete choice paradigms in which even mean-zero noise can systematically push people in the direction of seeming non-standard preferences. To illustrate, if a researcher offers a series of binary choices between a lottery and a safe payment where the lottery always has a higher expected value, then for a large class of utility functions noise will systematically push in the direction of seeming small-stakes risk aversion. In this spirit, recent experiments have shown that noise can spuriously generate or exaggerate a seeming demand for commitment in intertemporal decision-making (Carrera et al., 2022), seeming time inconsistency (Strack and Taubinsky, 2021; Chakraborty et al., 2017), seeming preferences for information avoidance (Exley and Kessler, 2021), seeming common-ratio Allais violations (McGranaghan et al., 2022), a seeming preference for certainty (Vieider, 2018), exaggerated estimates of risk aversion or risk lovingness (Enke and Shubatt, 2023; Belzil and Jagelka, 2020), exaggerated estimates of prosociality (Bao and Pei, 2023), seeming violations of axioms in general (Nielsen and Rehbeck, 2022), and spurious linkages between demographics and preferences (Andersson et al., 2016; Gillen et al., 2019). For example, half of the people who take up commitment contracts for higher gym attendance also take up commitment contracts for lower gym attendance, a reflection of noisy decision making (Carrera et al., 2022). Similarly, in the moral wiggle room literature, half of the people who exhibit “strategic ignorance” about their action’s payoff consequences for another person also exhibit “strategic ignorance” when their actions only concern their own payoffs (Exley and Kessler, 2021).

These results do not (and are often not meant to) show that the original behavioral motivations do not exist. However, they do show that (i) some of the evidence put forward in favor of non-standard utility can be confounded by noise and (ii) that we have probably

overstated their quantitative importance.

4 Comparative Thinking

As emphasized in Section 2, navigating tradeoffs across dimensions requires intensive information processing, both because it involves an element of aggregation and because it requires serious introspection about how to bring multiple problem dimensions into a “common currency.” A second simplification strategy to reduce the difficulty and costliness of such information processing is to primarily rely on comparative thinking: within-dimension comparisons rather than aggregation across dimensions.

The insight that people’s assessments of value, utility, quality and the like are typically relative rather than absolute in nature is one of the most prominent stylized facts in behavioral economics. For example, the basic idea that people implicitly or explicitly compare relevant quantities with (normatively irrelevant) comparison points or reference points features prominently in theories such as prospect theory and its variants (Kahneman and Tversky, 1979; Kőszegi and Rabin, 2006), disappointment and regret theory (Loomes and Sugden, 1982; Gul, 1991) and models of habit formation (Rozen, 2010). Even recent models of attention and salience (Bordalo et al., 2022c), stereotyping (Bordalo et al., 2016) and caution (Cerrei-Vioglio et al., 2022) are implicitly models of comparative thinking.

A leading interpretation of reference point effects in choice is that people actually have reference-dependent utility functions (in terms of their deep wants). I here summarize research suggesting that comparative thinking, instead, reflects a simplification strategy to deal with the difficulty of information processing across dimensions.⁶ The simple intuition is that it is often cognitively difficult to produce an absolute assessment or valuation (e.g., how much is this wine worth to me?), yet relatively easy to produce a relative assessment (e.g., this wine is better than the one I had yesterday). The general idea that reference points and comparison effects are linked to the cognitive difficulty of across-dimension aggregation is articulated in a range of decision theoretic contributions (e.g., Masatlioglu and Ok, 2005; Ok et al., 2015; Cerrei-Vioglio et al., 2022).

I will review three types of comparative thinking that have been discussed in the literature: (i) pairwise comparison effects in choice; (ii) comparisons to normatively irrelevant quantities (reference point effects); and (iii) sequential contrast effects in quality assessments. Across these categories, the overarching principle will always be the cognitive difficulty of aggregating different problem dimensions.

⁶From this perspective, comparative thinking is related to narrow bracketing (Read et al., 1999; Rabin and Weizsäcker, 2009; Ellis and Freeman, 2020; Lian, 2021). Narrow bracketing says that, when a decision maker encounters two decision problems, s/he considers them separately rather than integrating them with each other. Comparative thinking says that, when a decision maker encounters a single decision problem with multiple dimensions, s/he considers each dimension (partly) separately rather than first integrating across dimensions.

The literature on the cognitive foundations of comparative thinking is one of the least-developed ones that I cover in this review. Some of the conclusions I suggest below are more tentative than the ones drawn in other sections.

4.1 Comparison and Reference Point Effects in Choice

4.1.1 Pairwise Comparisons

To start, consider choice problems that are colloquially referred to as “simple” – choice sets comprising two goods with two dimensions each, such as choosing between an expensive high-quality wine and a cheap low-quality wine, or between a high probability chance to win a small amount and a small probability chance to win a large amount. The literature has distinguished between two different types of mental aggregation processes: (i) comparisons of aggregations – the decision maker first aggregates the components within each option and then compares the aggregated values across options; and (ii) aggregations of comparisons – the decision maker first compares the options component-by-component and then aggregates these comparisons (Tversky, 1969). Under neoclassical theory, these two processes are generally equivalent but in an emerging class of behavioral models they are not (e.g., Lanzani, 2022a).

In practice, many choice processes appear to reflect more closely comparison processes of type (ii), in particular when the decision is more difficult. Arieli et al. (2011) study this question using eye-tracking. They find that a large fraction of eye movements suggests within-dimension comparisons. Most relevant from the perspective of this review’s emphasis on the difficulty of aggregating tradeoffs, they find that the frequency of within-dimension comparisons sharply increases when the complexity of the lotteries is increased.⁷

The general idea that people engage in within-dimension comparisons rather than first aggregating within each option is also supported by a recent stream of work that empirically highlights that the difficulty of comparing options (within dimensions) has large effects on choice noise and self-reported uncertainty (Enke and Shubatt, 2023; Shubatt and Yang, 2024).

4.1.2 Comparisons to Normatively Irrelevant Reference Quantities

In principle, engaging in pairwise within-dimension comparisons need not reflect information-processing limitations because, after all, the decision maker compares two relevant quan-

⁷This evidence is linked to earlier research on both preference reversals and so-called joint-versus-separate evaluation effects (Hsee, 1996; Hsee et al., 1999). This literature documents that people’s evaluation of choice options typically differs between contexts in which only one option is presented and one in which people are asked to produce evaluations of two options. Because people are asked to make the same decision in both contexts, the leading explanation for a difference in evaluations is that people engage in a pairwise comparison process when two options are shown. Indeed, Hsee (1996) provides direct evidence that joint-versus-separate evaluation effects vary systematically with how difficult it is to value the separate product dimensions.

tities with each other. In a second category of comparative thinking effects, on the other hand, people compare with normatively irrelevant quantities. This is the typical domain of reference point effects, where people compare utility-relevant outcomes with comparison points such as the status quo, what they usually get, what they expect to get, and so on.

Recent work suggests that comparisons to reference points likewise reflect information-processing constraints. Just like in Section 3.2, this body of work can be partitioned into two distinct – but conceptually related – explanations: (i) the difficulty of aggregating tradeoffs across dimensions (or bringing different dimensions into a common currency) and (ii) efficient coding. Both of these accounts explain diminishing sensitivity away from a reference point as resulting from information-processing imperfections, though for different reasons. Accounts of tradeoff complexity rely on the idea that people’s decisions are very sensitive around reference points because they perfectly understand (in an ordinal sense) whether an outcome is better or worse than a comparison point, but not necessarily by how much. Accounts of efficient coding rely on the idea that people’s decisions are very sensitive around reference points because these are points with which they have much experience and values are, hence, encoded with less noise.

Difficult tradeoffs and reference point effects. To illustrate, suppose that a consumer is choosing between an expensive high-quality wine and a cheap low-quality wine. Suppose that the consumer is maximally uncertain about their exchange rate between money and quality of wine and can, hence, only produce *ordinal* assessments: whether or not one attribute is better than another one. For example, if the consumer has a reference point of medium quality and very high price (even higher than the expensive wine), s/he may opt for the high-quality wine because it delivers a “gain” in both dimensions relative to the comparison point, whereas the cheap low-quality wine would only produce one positive comparison. In this example, the consumer exhibits reference-dependent behavior purely for cognitive reasons, because s/he cannot aggregate across problem dimensions.

Enke and Graeber (2024) document that reference point effects indeed strongly vary with the difficulty of across-dimension aggregation. When the difficulty of bringing different problem dimensions into a common currency is exogenously decreased, classical reference point effects become considerably weaker or even disappear. For example, the widely-studied dependence of effort supply on past wages (e.g., Camerer et al., 1997; Abeler et al., 2011) almost entirely disappears when the difficulty of trading off money and leisure is experimentally reduced.

Efficient coding. Woodford (2012a,b) and Frydman and Jin (2022) suggest that reference point effects – in particular diminishing sensitivity away from reference points – reflect efficient coding of quantities. The idea is that the brain optimally encodes quantities in a more precise way (and thus produces a higher sensitivity of decisions) when these quantities

are encountered more frequently. According to this logic, reasonably large gains and losses are mentally compressed towards zero because they are less frequently encountered. To provide evidence for these ideas Frydman and Jin (2022) document that the sensitivity of decisions (and, hence, of prospect theory's implied value function) strongly increases in how much experience a decision maker has with the relevant quantities.

The idea of efficient coding naturally generates diminishing sensitivity around frequently-encountered points. Yet efficient coders do not *compare* quantities to a reference point. Thus from a psychological perspective efficient coding is an unlikely explanation of comparative thinking (though it likely contributes to diminishing sensitivity).

Normalizations. The idea that limits on information-processing capacity drive comparison effects implicitly also appears in work on relative thinking and normalization. This literature effectively says that any given within-dimension difference between two choice options matters more for decisions the smaller other “comparison quantities” in that dimension are. Different models take different views of what this “comparison quantity” is. Models of divisive normalization posit that, when people compare two options along one dimension, they normalize the difference by the average value in the set (Soltani et al., 2012; Webb et al., 2021; Landry and Webb, 2021). Relatedly, the salience model of Bordalo et al. (2012, 2013, 2022c) also features an averages-based normalization (in addition to a “differences-stand-out” element that will be discussed in Section 5). In both classes of models, a given difference in attribute values matters more for decisions when the average attribute is small. This captures the intuition that the difference between 30 and 20 seems larger than the difference between 530 and 520.

Models of relative thinking, on the other hand, posit that within-dimension comparisons are not normalized by the average but by the range of values in the set (Bushong et al., 2021). This captures the intuition that the difference between 30 and 20 seems larger when the third option is 19 than when it is 0.

Regardless of the specifics of the models, they all afford a straightforward interpretation through the lens of information processing, in particular the difficulty of navigating trade-offs across dimensions. When people do not really know how to translate a given problem dimension into a “common currency” (e.g., utils), then the range or magnitude of attribute values in the choice set provides a plausible cue about how large and important the attribute difference in that domain actually is in utility terms. From this perspective, the simplification strategy of weighing more heavily dimensions with small range or average attributes may “make sense”.

A range of recent experimental contributions has provided evidence in favor of comparison-based thinking in line with range- or average-based normalization. This includes settings such as lottery choice (Soltani et al., 2012; Frydman and Mormann, 2018), consumer choice (Somerville, 2022; Dertwinkel-Kalt et al., 2017) and effort provision (Bushong et al., 2021).

Across these contributions, the key message is that within-dimension differences loom larger (and, hence, matter more for decisions) when the average or range of values in the set is smaller.

Future work could helpfully study to what degree these range- or average-based normalization effects depend on the difficulty of aggregating tradeoffs or bringing different dimensions into a common currency.

4.2 Quality Assessments

One indication that reference point effects are ultimately cognitive (rather than preferences-based) in nature is that there is a large body of evidence on the existence of comparative thinking with respect to normatively irrelevant points also in decision domains that do not involve preferences. The leading example is the so-called sequential contrast effect, which holds that – in magnitude or quality estimation – a quantity appears larger if the previously-encountered quantity (the comparison point) is smaller. For example, economists have shown that assessments of candidate quality sharply decrease in the quality of the immediately-preceding candidate (Bhargava and Fisman, 2014; Radbruch and Schiprowski, 2022) or that positive earnings surprises have larger impacts on stock prices if yesterday's earnings surprise was negative (Hartzmark and Shue, 2018). The common rationale behind these patterns is that when people try to judge the quality or value of an object, they do not really know how to bring the different dimensions into the same currency. For instance, in hiring contexts, it is cognitively easier to know that “candidate A has better math grades than candidate B” than to produce a precise quantitative evaluation of how candidate A's math grades should map into an overall quality assessment on a scale from, say, zero to ten. Enke and Graeber (2024) provide experimental evidence for this. They document that when the difficulty of translating different dimensions of quality into a common currency is exogenously increased, sequential contrast effects become considerably stronger.

Summary. In summary, the entirety of the evidence summarized above suggests that the process of comparing within dimensions is a simplification strategy used to address the cognitive difficulty of translating different dimensions into a common currency. As noted above, the literature on the information-processing origins of comparative thinking and reference point effects is less developed than some of the other literatures reviewed here. This is true both empirically and in terms of theory. For instance, theoretical models that formalize comparative thinking as resulting from information-processing imperfections are still relatively rare at this point.

5 Reducing Cardinality

The simplification strategies summarized so far rely on the idea that people use all of the available information, albeit in a noisy and comparative way. According to this literature, people make mistakes, but they never fundamentally misunderstand a problem or even entirely ignore an aspect that actually matters. These simplification strategies predominantly apply to situations in which the problem is reasonably low-dimensional, such that all variables and options can actually be considered, compared and approximated. Yet in many economic applications, the set of variables to be considered is large. In these contexts, it is implausible that the decision maker can engage in a comparison process within each relevant problem dimension, or noisily approximate the value of each option. Instead, a more plausible simplification strategy is to work with a subset of the available information (or a subset of variables), and to decide based on what's top of mind due to selective attention and / or selective memory.

At a high level, the literature on reducing the dimensionality of the problem through selective attention and memory can be partitioned into two streams. First, work that explicitly posits an optimizing or goal-driven element: people mostly attend to, and remember, those aspects that are more important for the problem at hand. This is the literature on rational inattention, sparsity and optimal bounded memory. Second, work that emphasizes the importance of stimuli and context: people mostly attend to, and remember, those aspects that are cued by the decision environment, regardless of whether or not they are actually important for the problem at hand.

While these two approaches appear to reflect fundamentally different approaches to thinking about simplification strategies (and while they certainly are very different from each other at a formal level), I believe that many of the stimulus-driven simplification strategies are likewise best understood as broadly-optimal solutions to the basic constraints of limited attention and limited memory. For example, models of stimulus-driven attention and salience posit that people predominantly attend to problem dimensions that exhibit large variation (e.g. Bordalo et al., 2022c; Kőszegi and Szeidl, 2013). While this may be wrong in any given problem, it appears that such a strategy does make sense “on average” (across problems), and indeed optimizing models of attention allocation generate exactly this prediction (Gabaix, 2019; Maćkowiak et al., 2023). Likewise, models of stimulus-driven recall emphasize that people predominantly recall information that was experienced in a context that is similar to the current decision context (Mullainathan, 2002; Bordalo et al., 2023b). Again, this may lead to mistakes in any given situation but appears to broadly make sense “on average”. After all, situations encountered in the past that are similar to the current decision are likely more informative than completely unrelated decisions.

A useful way to think about optimizing versus stimulus-driven approaches to attention and memory is, hence, that both capture ecologically-rational heuristics, albeit of the form

that in optimizing approaches the attention and memory allocation is optimal even *for the specific problem at hand*, while in the stimulus-based approaches it is only optimal *on average*.

In what follows, I first discuss evidence of situations in which people entirely ignore or heavily underweight some aspect of the problem, and then summarize models that formalize psychological forces that can generate such effects.

5.1 Evidence on Incomplete Representations

By “incomplete representations” I mean situations in which people’s decisions suggest that some problem aspect doesn’t even come to mind and is fully neglected, for example when some people entirely miss an edge in the network that represents the true causal structure of the economy. As a result, decisions tend to exhibit a pronounced multi-modal structure that is typically absent when noisy approximations or comparative thinking dominate the decision process.

A useful way to descriptively organize these patterns – through not to adequately represent the cognitive mechanisms at work – is through directed acyclic graphs (Spiegler, 2016, 2020). We can distinguish between two classes of regularities that the literature has accumulated: (i) people entirely miss an edge in the deeper structure of the network (or “system”) that produces the data, and excessively focus on what’s most immediate or visible; and (ii) people ignore even some of the immediately presented evidence.

System neglect. There is much evidence in psychology and decision research to suggest that people tend to excessively focus on the salient, immediately visible “output” (data) that a causal structure produces, rather than on the features of the underlying data-generating process (e.g., Massey and Wu, 2005; Fiedler and Juslin, 2006). The majority of experimental results that have this structure are found in the domain of belief updating. Consistent with the emphasis in this review, all of these experiments are characterized by their focus on how people aggregate multiple pieces of information. A repeated pattern in the literature is the existence of two “types” that can be understood as having different mental models of the data-generating process: those who are rational and those who simplify by entirely ignoring a link and implicitly treating the observed messages as if they equaled the underlying signals. This simple logic of an incorrect mental model (that equates signals and messages) descriptively organizes the strongly bi-modal response patterns in experiments on correlation neglect (Enke and Zimmermann, 2019), selection neglect (Enke, 2020; Jin et al., 2021; Barron et al., 2023), omitted variable bias (Graeber, 2022) and de-Groot updating in social networks experiments (Grimm and Mengel, 2014; Chandrasekhar et al., 2020). A recurring result in this literature is that drawing people’s attention to the previously-neglected aspect of the data-generating process has large effects on behavior, consistent with the idea that incorrect mental representations are partly driven by selective attention.

Closely related to this evidence on system neglect is work on neglect of indirect effects, most notably of equilibrium effects (Dal Bó et al., 2018; Andre et al., 2023b). This work suggests that people do not pay sufficient attention to how others' behavior is driven by fundamentals, and even when they do, they pay insufficient attention to how such behavioral responses would impact aggregates. Again, we can think of these patterns through the lens of a simplified representation of the data-generating process in which certain edges are eliminated. This work on incorrect representations of data-generating processes is increasingly used also in applied contexts, including mental models of the stock market, inflation, and the macroeconomy (Andre et al., 2022, 2023a,b).

Data neglect. The general idea that people simplify by entirely ignoring certain problem aspects is also present in work on belief formation that studies how people combine base rates and likelihoods in experimental two-states-binary-signals belief updating paradigms. This is again a problem of aggregation, except that now base rate and likelihood need to be traded off against each other, rather than different signals. A recurring pattern in the literature is that, in these problems, reported posterior beliefs also often exhibit a multi-modal structure: (i) some people who are rational; (ii) some people who report back the base rate; and (iii) people who report back the likelihood (Fan et al., 2023; Bordalo et al., 2023a; Esponda et al., 2020). This suggests that people selectively truncate problem aspects. Bordalo et al. (2023a) present a model that explains these patterns as resulting from selective attention – people are modeled as paying attention to only one statistic (the base rate or the likelihood ratio) depending on which one is more salient.

Similarly, Conlon (2023) documents that people tend to entirely ignore some dimensions in multi-attribute choice. As in many other contributions in this literature, nudging people's attention to certain aspects of the problem has large effects on observed behavior.

In the field, Hanna et al. (2014) study the behavior of seaweed farmers who need to take into account many input dimensions. These farmers entirely neglect the relevance of a particularly important input feature even though data on it is readily available. Consistent with an incorrect mental model driving a failure to optimize, the farmer's decisions substantially improve when their attention is drawn to previously-unnoticed relationships. Just like in the case of the lab experimental evidence, there is hence reason to believe that selective attention is a major driver of incorrect mental models.

Contingent reasoning. Many economic decisions require people to reason through future (unrealized) hypotheticals. Yet the need to aggregate multiple future hypotheticals potentially introduces severe cognitive information-processing demands, for two reasons. First, people do not necessarily organically mentally represent the full state space in the ways our models envision. Second, even conditional on representing the state space in a certain way, information processing is required to aggregate across different states, some of which

may be more top of mind than others.

A rapidly-growing recent literature documents that people often mentally misrepresent decision problems through failures of contingent reasoning, by failing to properly attend to the decision-relevant states of the world. This literature, reviewed in Niederle and Vespa (2023), grew out of the experimental game theory literature, in particular the work on overbidding in auctions and cursed equilibrium (Kagel, 1990; Eyster and Rabin, 2005).

Recent work on this topic is a prime example of research on cognitive foundations in its desire to document that the failure to properly condition on (and attend to) future hypotheticals is a driver behind multiple behavioral anomalies. This idea was first articulated by Esponda and Vespa (2023), who document that overbidding in auctions, mistakes in voting and Ellsberg paradoxes are all partly driven by a failure of contingent reasoning. A main empirical innovation in this work is to switch the need to reason through multiple hypotheticals on and off, thereby directing people's attention to the main decision-relevant state of the world. For example, Esponda and Vespa (2023) document that the two-urn Ellsberg paradox is significantly weakened once people are prodded to focus on the decision-relevant state, in much the same way as common-value voting improves once people are induced to focus on the state of being pivotal (Esponda and Vespa, 2014).

In conceptual terms, one way to understand this literature is that people find it cognitively challenging to aggregate across many unrealized states of the world, and thus form a mental representation of the problem according to which all states are "equally relevant". In this sense, work on contingent reasoning is a prime example of the importance of aggregating tradeoffs across dimensions (states) that I highlight in this review.

The broad idea of failures in contingent reasoning driving behavior has been leveraged to understand behavior across multiple different contexts, including the winner's curse (Charness and Levin, 2009; Martínez-Marquina et al., 2019; Nagel et al., 2023), common-values voting (Esponda and Vespa, 2014), mistakes in annuity take-up (Luttmer et al., 2023), public goods provision and redistributive behavior that is insensitive to circumstances (Andre, 2022; Calford and Cason, 2022), learning from market prices (Ngangoué and Weizsäcker, 2021) and mechanism design (Kendall and Chakraborty, 2022).

This long list indicates significant progress in the quest for cognitive foundations that generate multiple anomalies. After all, each of the various literatures mentioned above had developed idiosyncratic explanations and models for the phenomena they sought to explain. Yet taking a step back and looking across literatures, recent work suggests that many anomalies reflect that people entertain mental representations in which the decision-relevant state is not top of mind.

Building on this body of evidence, theorists have developed models that formalize people's difficulty in thinking through hypotheticals and draw out corresponding implications for strategic decision making (e.g., Li, 2017; Cohen and Li, 2022; Pycia and Troyan, 2023). See the discussion and review in Li (2024).

5.2 Goal-Directed Attention and Memory

A large and growing literature studies the role of goal-directed attention and memory. The main idea is that people attend to, and remember, what they believe to be important. Again, a main motivation for this literature is the existence of information processing constraints – too many things go on in our lives, too many variables need to be attended to and too much information needs to be stored in memory.

Rational expectations approaches. The largest and most prominent literature on directed attention is the one on rational inattention and sparsity (Gabaix, 2014). This refers to a model of information acquisition in which the amount of information that is acquired is subject to a cost, often parameterized by the expected reduction in entropy that information induces (thus more informative signals are costlier). The theory literature on goal-driven inattention is reviewed in great detail in Maćkowiak et al. (2023) and Gabaix (2019), so I here focus on the extant experimental evidence.

Direct tests of the rational inattention model often rely on so-called psychometric tasks from psychology. In these tasks, subjects solve simple perceptual decision problems such as determining whether the majority of colored dots on a screen is red or blue (Caplin and Dean, 2015; Caplin et al., 2020; Dean and Neligh, 2023). The researchers then vary aspects of these decision tasks to study how performance varies as a function of incentives and problem setup. These experiments reveal that rational inattention explains information search and decisions well in a qualitative sense (Caplin and Dean, 2015), with the important caveat discussed in Section 3 – that performance tends to be considerably less sensitive to incentives than predicted by canonical rational inattention models (Dean and Neligh, 2023).

At this point, most direct tests of rational inattention rely on perceptual decision tasks. A fruitful path forward would be to study whether behavior in more standard economic choice tasks (even if in the lab) can be understood through the lens of the rational inattention framework. A notable exception is Bartoš et al. (2016) discussed in Section 10.

Model-driven inattention: Learning what matters. In models of rational inattention and sparsity, the decision maker has rational expectations about which variables matter how much. A second class of models instead assumes that the decision maker needs to learn the structure of the problem over time, in particular which variables influence the outcome of interest in which way (Schwartzstein, 2014; Gagnon-Bartsch et al., 2021; Fudenberg and Lanzani, 2023). In these models, people attend to variables or dimensions that they expect to be relevant given their current worldview, which need not necessarily be correct. As a result, people may persistently fail to learn important patterns in the data if their initial incorrect model says that the corresponding variables are unimportant.

The fact that this class of models does not rely on optimal cognition links it to procedural

models of directed attention such as Gabaix and Laibson (2000), Gabaix and Laibson (2005) and Gabaix et al. (2011). These papers propose and experimentally test cognitive algorithms for information acquisition. For example, in searching for information, people may behave as if the current search step were the last one (Gabaix et al., 2011) or they may ignore low-probability events (Gabaix and Laibson, 2000).

A notable feature of this class of models is that – unlike most of the theoretical approaches and empirical regularities discussed earlier – they produce discretely misspecified mental models. In other words, these models capture the idea that when a problem is difficult, one way to simplify is to “truncate” – to entirely ignore aspects of the problem that are difficult to understand or deemed unimportant.⁸ As a result, models in this vein (and corresponding experimental evidence) often generate large spikes at incorrect decisions, rather than the usually more continuous mistakes described by models of behavioral attenuation or what’s top of mind.

There is direct evidence for the idea that an initial misrepresentation of a decision problem affects subsequent attention allocation, in a way that an initial misrepresentation can induce a permanent failure to learn. Esponda et al. (2020) present experiments in which people receive stochastic feedback about their performance in a base-rate neglect problem. They find that when people are initially given a chance to form a mental representation of the problem, they do not adjust their beliefs in response to feedback (presumably because they believe they have understood the problem and don’t need to pay attention to the feedback). In contrast, when people are not given an opportunity to come up with a mental problem representation because they do not know the precise problem primitives, feedback strongly affects reported beliefs. This suggests that excessive confidence in an initial incorrect model can persist because it induces people to stop paying attention to feedback.

Model selection. A final example of the general idea that people simplify excessively complex problems by entirely ignoring some problem aspects is the literature on model selection and model-based persuasion. Consider a decision maker who entertains multiple subjective mental models of the world. For example, the decision maker is looking to update his beliefs about future inflation, but may be uncertain over whether oil price shocks do or do not impact inflation. In situations like these, a Bayesian would compute beliefs conditional on each model of the world and then average across potential models. This is again a problem of aggregation that could pose a high degree of representational complexity if there are many potential models.

A recent theory literature posits that people simplify this problem by attending to only one model and updating beliefs purely based on this one. For example, people may select the

⁸Theorists have devoted much attention to formalizing, characterizing and investigating the dynamic properties of generic incorrect mental models, without taking a stance on what the concrete misspecification is (e.g., Spiegler, 2016, 2020; Esponda and Pouzo, 2016; Bohren and Hauser, 2021; Heidhues et al., 2018; Fudenberg et al., 2022; Lanzani, 2022b).

model that offers the best fit of the available data given people's prior beliefs (Schwartzstein and Sunderam, 2021, 2022; Aina, 2021). This is an example of goal-driven attention just like the ones above because people are assumed to (exclusively) attend to worldviews and narratives that are compelling. Barron and Fries (2023) provide experimental evidence that is consistent with such models. They document in a sender-receiver context that receivers are more likely to adopt the mental model that a sender with strategic motives suggested to them if that model has a good fit. In contrast, Kendall and Charles (2022) find that their subjects engage in model-averaging behavior, akin to the way a Bayesian averages models.

Goal-directed memory. A nascent literature studies optimal memory when there is a cost or another type of hard constraint on what can be remembered. In contrast to the literature on similarity-based recall, this body of work derives features of recall based on what's optimal (subject to a constraint) rather than by formalizing known recall biases from psychology. Examples in this literature include the contributions of Wilson (2014), Azeredo da Silveira and Woodford (2019); Da Silveira et al. (2020) and Bakhtin et al. (2023). Each of these papers focuses on laying out how (optimal) costly memory can explain a variety of biases in information processing or forecasting. At this point, however, direct empirical evidence on this class of models is missing.

5.3 Stimulus-Driven Attention

The literature on stimulus-driven attention rests on the idea that people intuitively overweight aspects of the problem that are salient or stand out. A large literature in psychology documents that low-level cues (such as red font) attract disproportionate attention. A main objective in the economics literature is to formalize "economic" determinants of stimulus-driven attention, such as how the distribution of prices (rather than how the prices are displayed) shapes choice.

Two interrelated features characterize this literature. First, work on the topic has heavily relied on models of context-dependence. In choice contexts, this means that a decision maker's evaluation of a given choice option partly depends on the other options in the set. In inference and prediction contexts, it means that the decision maker's assessment of a given hypothesis partly depends on which alternative hypothesis is considered. Second, and relatedly, this literature implicitly or explicitly emphasizes that differences (or variability) attract attention. For example, a price may stand out because it appears unusually large or small, or an attribute may stand out because it is much more frequent in one population than in another. Because such differences attract attention, they receive disproportionate weight in the decision process. As a result, in contrast to the work on behavioral attenuation and comparative thinking summarized above, research on what's top of mind generates that people's relative weights between different problem components are systematically distorted.

5.3.1 Choice

To illustrate, suppose that a decision-maker’s optimal decision (i.e., their utility-maximizing choice or the normatively correct posterior belief) is given by a weighted average of different problem components, $a^*(\theta) = \sum_j \beta_j \theta_j$. Then, much of the literature can be summarized as representing special cases of attention-distorted decision rules of the form:

$$a = \sum_j w(\underbrace{\beta_j}_{\text{Normative weight}}, \underbrace{\theta_j, \theta_{-j}, c}_{\text{Context dependence}}) \theta_j, \quad (5)$$

where $w(\cdot)$ is an attention-based weighting function that captures “what’s top of mind” or “what stands out”. Importantly, this weighting function potentially depends on the attribute values of all options in the set, θ_i and θ_{-i} , as well as contextual cues, c . This generates context-dependence because the choice between two options can be affected by the characteristics of other options through its impact on the weighting function. Similarly, the assessment of one hypothesis can be affected by the characteristics of other hypothesis that are actively considered.

The models and empirical applications in the literature differ in how they specify the weighting function, as discussed below. However, the overarching principle is always the same – that a dimension or problem component receives disproportionate weight when it attracts attention, which is the case when differences (or relative differences) are large.

A final general comment concerns the mathematical formalism that these literatures often rely on. Unlike most of the models reviewed above, they do not explicitly model information processing (as in, for example, when a decision maker receives a noisy signal about what their optimal policy function looks like). Instead, as illustrated in equation (5), they put a reduced-form weighting function into the utility function. However, the interpretation of this is usually not that people’s experienced utility actually depends more strongly on those dimensions that receive high decision weights but, instead, that people selectively allocate attention due to bounds in information processing capacity. In other words, these are usually mathematical representations of perceived (“decision”) utility rather than experienced utility.

Consider the salience model proposed and reviewed in Bordalo et al. (2012, 2013, 2022c). In this model, the weight of dimension j for good i increases in the distance between i ’s attribute value and the average attribute in dimension j , normalized by the average attribute in the dimension, $w_i^j = f\left(\frac{|\theta_i^j - \bar{\theta}^j|}{\bar{\theta}^j}\right)$. Thus in this model, percentage differences are said to attract attention, in a way that follows an explicit comparison (or reference point) logic.⁹ For instance, in a two-item menu, the relative weight of each dimension is pinned down by a pairwise within-dimension comparison, such that the dimension receives higher weight if

⁹Various papers have discussed the connection between salience and earlier implicit reference points models such as regret theory, see Herweg and Müller (2021); Ellis and Masatlioglu (2019); Lanzani (2022a).

the options are more spread out in that dimension (in percentage terms).

The focusing model of Kőszegi and Szeidl (2013) instead assumes that the weight of each dimension is given by the (utility-weighted) range of attribute values in that dimension, $w_j = f(\max u(\theta_i^j) - \min u(\theta_i^j))$. Similarly to the salience model, this captures the idea that differences and variation attract attention. One interpretation of these models is that they capture a quasi-rational heuristic of paying special attention to dimensions in which the available options exhibit lots of variation (clearly a dimension can be ignored if all options in the set are equal in that dimension). This model can, hence, be understood as a heuristic version of an optimizing attention model.

As summarized in Bordalo et al. (2022c), a wide range of empirical evidence can be interpreted through the lens of context-dependent choice models, including decoy effects, context-dependent willingness-to-pay and demand for positively skewed financial assets. Some of the more direct experimental tests of specific elements of context-dependence include Dertwinkel-Kalt and Köster (2020a,b) and Bruhin et al. (2022).

Interestingly, the models of Kőszegi and Szeidl (2013) and Bordalo et al. (2013) are very similar at a technical level to the relative thinking (normalization) model of Bushong et al. (2021), except that the latter assumes that – building on the psychology of comparison and normalization – a higher range of attribute values *decreases* the decision weight that's attached to a particular dimension. Perhaps the leading view in the literature is that these models – while mathematically very similar – capture two different types of psychology that play out in different contexts. The “differences-attract-attention” view is probably correct in high-dimensional problems in which attention is a scarce resource. For example, when deciding between 10 options, each of which has 10 different dimensions, it is plausible that people focus on goods and dimensions that stand out because they exhibit large variation. In contrast, the relative thinking logic applies to smaller choice sets in which people can directly compare all relevant options and dimensions.

The extant experimental evidence is consistent with this view. First, as noted above, experiments in low-dimensional contexts tend to find that the relative thinking and comparison logic dominates (Soltani et al., 2012; Somerville, 2022; Bushong et al., 2021; Dertwinkel-Kalt et al., 2017; Frydman and Mormann, 2018). These are all experiments in which subjects choose among two or three goods, each of which has a very small number of dimensions (typically two). In contrast, more high-dimensional experiments tend to find support for the differences-attract-attention perspective. For example, Dertwinkel-Kalt et al. (2022) document that intertemporal decisions over many future dates are heavily biased towards large differences. Similarly, Bohren et al. (2024) ask experimental participants to choose between two assets with a large number of distinct payout states, and find that people heavily overweight those states in which the two assets have large payout differences.

The literature on field applications of inattention is relatively sizable, in particular in the context of shrouding of attributes (e.g., Chetty et al., 2009; Brown et al., 2010; Bradley

and Feldman, 2020). However, few papers test the more specific predictions that different models of salience or focusing make. An exception is Hastings and Shapiro (2013) who show that salience effects explain a part of observed consumption changes when gasoline prices rise.

Visual salience. Most of the experimental literature on salience and focusing investigates the predictions of models in which context-dependence is driven by payoffs or objective attribute features (what one might call “economic salience”). A complementary line of work has highlighted that attention allocation also depends on visual salience in ways that are orthogonal to economic salience. For example, Li and Camerer (2022) document in a consumer choice experiment that participants are less likely to make choice mistakes if the higher-value option is also the visually salient one. Notably, the setup of their experiment is such that an objectively correct solution exists, such that the only relevant cognitive difficulty consists of mentally aggregating the different constituent components of a choice options. This highlights that responding to visual salience is a simplification strategy to the difficulty or costs of aggregating and trading off problem dimensions, much like behavioral attenuation is.

In a related study, Bose et al. (2022) show that forecasts of autocorrelated processes based on standard price charts significantly depend on which part of the price chart is visually salient. For example, when positive past returns patterns are visually salient, people overestimate future returns and invest more into the respective stock. These results are derived by importing machine learning algorithms that predict the visual salience of different parts of images.

5.3.2 Beliefs

Closely related to work on salience and focusing in choice is research on inference and prediction that highlights the importance of the differences-attract-attention principle. According to a growing body of work, people disproportionately overweight information that is representative of the underlying population. Here, representativeness is again defined in a comparative way (e.g., Gennaioli and Shleifer, 2010; Bordalo et al., 2016, 2018, 2022b). In a nutshell, the decision maker asks: how likely is it that I would observe this information under hypothesis (or state of the world) A, relative to how likely it is under hypothesis B? If this difference is large, people are said to overweight the information relative to their prior, and hence overreact.

Notice the close analogy between models of differences-attract-attention in choice and beliefs – in both contexts, the models and empirical evidence suggest that people respond to real differences, yet exaggerate these differences when they are large (because they attract attention).

While the idea of representativeness goes back decades (e.g. Kahneman and Tversky, 1972; Camerer, 1987), it was reinvigorated with the cognitive turn, partly because it appears to reflect the working of attention and memory (Bordalo et al., 2021, 2022b, 2023b; Wachter and Kahana, 2023).

Evidence. Various lab experiments have confirmed that reasoning in the form of representativeness generally produces a type of overreaction effect. For example, representativeness can be understood as underlying stereotyping in the domains of gender (Bordalo et al., 2019), politics (Bordalo et al., 2016) and information processing (Esponda et al., 2023). In all of these contexts, the main empirical pattern is that when a true difference between groups or states of the world exists, people exaggerate it.

The general idea that people's beliefs reflect the exaggeration of true differences has also been applied in ecological applications. Conlon and Patel (2022) document that undergraduate students strongly stereotype the link between college major choice and subsequent occupational outcomes – they exaggerate the likelihood that enrolling in a major will lead them to take up a job that is most representative of that major.

Complexity, what's top of mind, and behavioral attenuation. The perspective of this review is that the tendency to overweight what stands out is a simplification strategy just like behavioral attenuation and comparative thinking (indeed, as noted above, focusing on what stands out constitutes a form of comparative thinking itself). This said, at this point, there are fewer contributions that have directly studied the link between imperfect information processing (or complexity) and acting based on what's top of mind. Two notable exceptions are Dertwinkel-Kalt and Köster (2020b) and Ba et al. (2022). Ba et al. (2022) propose a model in which the distinction between representational and computational complexity takes center stage. In a first step, the decision maker is assumed to reduce representational complexity by over-attending to states of the world that stand out (according to what is representative). In a second step, the decision maker reduces computational complexity through noisy approximation and resulting behavioral attenuation. Theirs is the only framework at this point that combines multiple different simplification strategies. They document experimental support for the model's predictions, in particular by showing that increasing representational complexity increases the importance of representativeness. In contrast, Dertwinkel-Kalt and Köster (2020b) use a lab experiment to document that people react to salient aspects of a portfolio choice problems regardless of whether or not it is computationally complex.

5.4 Similarity-Based Recall

Closely related to work on stimulus-driven (bottom-up) attention is research on stimulus-driven memory. Indeed, what's top of mind in the moment of making a decision is not only determined by the differences-attract-attention principle but also by which particular memories get cued by a decision context. To illustrate, a retail investor who struggles with the cognitive difficulty of processing and aggregating all available information about the market may only have specific instances from the past top of mind, based on which contextual cues he is exposed to today.

This raises the general question of which problem aspects tend to be top of mind because they get cued in memory. Probably the most important and most robust result in memory research is that recall is associative in nature (or similarity-based): people are more likely to remember things that are similar to what they observe today. Hundreds, if not thousands, of papers in psychology and decision neuroscience rest on this principle (see, for example, Tversky, 1977; Kahana, 2012; Baddeley, 2013). Thus work on the topic either implicitly or explicitly features a similarity metric between current and past decision contexts. Importantly, similarity is not necessarily only defined on objectively relevant features of the problem, but potentially also on contextual cues (narratives, images etc.).

5.4.1 Beliefs

To illustrate this class of models, suppose that a decision maker's true utility-maximizing decision is given by $a^*(w, x) = \sum \beta_i x_i$, where x_i are realizations of the attributes. For example, product demand is a weighted average of product attributes, or the rational posterior belief is given by a weighted average of various information signals. Further suppose that the attributes are not necessarily visible (or top of mind) today, meaning that the decision maker must activate or retrieve them from memory. Today, the decision maker faces contextual cues c^T , while the attributes were initially experienced in context c^H . Then, a strand of the literature on similarity-based recall can be understood as representing special cases of the decision rule

$$a^o(w, \theta, c^T, c^H) = \sum S(c^T, c^H) \beta_j \theta_j, \quad (6)$$

where $S(\cdot)$ is a similarity function. Thus in these models an attribute only affects decisions to the degree that it was experienced in a context that is similar to today's context. As noted above, here the "context" can include both the attribute realizations and payoff-irrelevant contextual features that affect similarity judgments and recall.

Theories. The main starting point for all theories of expectation formation based on memory retrieval is that, in reality, people rarely have access to concrete statistics that summarize

for them the relevant probabilities. Instead, people retrieve past experiences from memory and combine them into probabilistic forecasts. As a result, unlike the vast majority of models of belief formation in the (behavioral) economics literature, memory-based theories model how people recall personal experiences, rather than how they use statistics such as base rates.

The first example of this class of models is Mullainathan (2002) who formalizes the idea of associative recall in the context of expectation formation. A key prediction of this model as well as most subsequent ones discussed below is that of memory-driven overreaction. To illustrate, suppose that attributes (or news) x_i are correlated with the occurrence of certain contextual cues. For example, good news about the stock market always appear with images of bulls and upward-sloping trend lines, while bad news are usually associated with images of bears. Then, if a decision-maker receives positive news (and the associated images) today, then similarity-based recall could induce them to asymmetrically remember only those past news (or attribute realizations) that are similar to today's news. As a result, the beliefs and behavior of the decision maker look like they overreact to news, purely as a result of how the news affect how the decision maker reconstructs prior knowledge. In a nutshell, models of associative recall often generate a form of overreaction to news because they encompass both a direct and an indirect effect of news – the direct effect is the standard process of belief updating following news, while the indirect one is that recent news may affect what people recall from the past, which then also affects beliefs.

Kőszegi et al. (2022) use a similar setup to study excessive confidence swings. They observe that if low confidence cues the recollection of negative news about the self, there exist multiple fragile confidence equilibria, in which excessively high or excessively low confidence and selective the memories that they trigger sustain each other. As a result, akin to Mullainathan (2002), recent news can trigger large belief movements because they have both a direct effect and an indirect effect through people's mood and the selective memories that it induces.

Malmendier and Veldkamp (2022) apply the idea of similarity to the domain of interpersonal communication. They model a decision-maker who is more inclined to accept and act on news that “resonates”, for example because the sender has similar objective characteristics as the receiver.

Bordalo et al. (2023b) and Bordalo et al. (2022a) develop and apply a model of how people form beliefs and expectations when the recollection of events is governed by two principles: similarity and interference. First, events are more likely to be recalled when they are more similar to (or when they were experienced in a context that is more similar to) what is considered today. Second, events are less likely to be recalled the more similar they are to alternative hypotheses or events. While the principle of similarity is intuitively obvious, interference captures the idea – widely-studied in memory research – that human memory tends to “mix up” experiences and hypotheses when they are spuriously similar to each other.

To illustrate, consider the context of economic statistics (such as GDP growth or unemployment rates), which are routinely represented through numbers. Because people encounter numbers and statistics all the time, it is very difficult to remember specific unemployment statistics from two years ago – intuitively, the brain mixes up the various numbers it has encountered over the years. Bordalo et al. (2023b) and Bordalo et al. (2022a) show that the ideas of similarity and interference can reconcile various biases in judgment of probabilities. A main idea in these papers is that normatively identical situations can trigger different beliefs, either because the environment provides different cues or because the decision-maker actively considers different hypotheses, both of which can trigger the retrieval of different memories.

One of the main applications of memory-based theories of belief formation has been to financial markets. There are probably two natural reasons for this. First, a long literature documents that personal experiences exert a strong influence on financial and macroeconomic expectations (Malmendier and Nagel, 2016; Malmendier et al., 2021; Malmendier and Wachter, 2021; Malmendier, 2021). Second, as discussed above, associative recall provides a natural rationale for the existence of short-run overreaction to news, which is often observed in economic expectations (Bordalo et al., 2020a).

Bodoh-Creed (2020) presents a model that is conceptually tightly linked to Kőszegi et al.'s 2022 model of excessive confidence swings. In Bodoh-Creed's model, recall of positive or negative news depends on the decision-maker's mood, which is governed by recent news (as these trigger changes in the state of a stock market). For example, following a series of positive dividend announcements, the decision-maker is in a good mood and selectively recalls mostly positive news from the past, which leads him to overreact in the short-run. This model generates excess volatility in beliefs and market movements. Wachter and Kahana (2019) present a similar model and use it to explain a variety of stock market puzzles, including personal experience effects and financial crises.

Evidence. Recent experimental work has confirmed the role of similarity and interference for the formation of beliefs (Bordalo et al., 2023b). In particular, this work highlights that the presence of associations (cues) in combination with memory constraints indeed produces predictable overreaction of expectations to recent news (Enke et al., 2024b). These patterns are especially pronounced when information is conveyed in the form of memorable narratives rather than numerical statistics because narratives decay more slowly in memory (Graeber et al., 2022).

A growing literature studies how similarity-based recall shapes expectation formation and decision making in the wild. Jiang et al. (2022) use surveys to document that positive market returns cue investors to retrieve episodes of rising markets. Moreover, recalled experiences explain a larger share of variation in stock market expectations than actual experiences. Relatedly, Bordalo et al. (2022a) use surveys to show how similarity of health-related

experiences shapes beliefs about Covid fatality rates.

Work in psychology. Economic models of similarity-based recall directly build on exemplar models that psychologists developed in the 1980s and 90s and refined in the early 2000s (Kahana, 2012). As with the case of Bayesian cognitive noise models, the potential criticisms that behavioral economists working on similarity-based memory have encountered are very similar to the criticisms that were leveraged in cognitive psychology three decades ago. Most importantly, the potential lack of discipline in which contextual features enter the similarity function was first pointed out by psychologists many years ago (Medin et al., 1993) – essentially, the argument is that more or less any behavior can be explained with arbitrary similarity functions, just like more or less any behavior can be explained with arbitrary utility functions.

5.4.2 Choice

Bordalo et al. (2020b) propose a model of choice in which the decision maker engages in comparative thinking. The main difference to the approaches reviewed in Sections 4 and 5.3 is that in their model comparative thinking occurs with respect to a reference point that is determined by similarity-based recall. In a nutshell, the decision maker retrieves similar choice problems from the past (as these are top of mind) and then compares quantities such as prices or qualities with what he encountered in similar problems in the past. Thus in this model people’s valuation for goods depends on the context in which they are consumed because different contexts bring top of mind different “normal” situations. This approach, hence, naturally speaks to choice instabilities and context effects.

At this point, direct empirical work on the idea of a similarity-based reference point is scarce. While there is much evidence for similarity-based recall and the importance of contextual cues in general (see Sections 5.4.1 and 6), direct empirical investigations of the comparison process that people implement when they compare current choice options with options retrieved from memory would be very valuable to have at this point.

6 Categorizations and Analogies

All of the approaches reviewed so far rest on the idea that people actually attempt to solve the decision problem that is being presented to them by integrating the relevant problem aspects with each other. A different strand of the literature instead posits that people learn from the past – they may not even have a causal representation of why which action “works”, but they simply repeat what worked yesterday. This literature, in turn, can be partitioned into two strands. First, work in which people learn from exactly the same problem (as in reinforcement learning). Second, work in which people draw analogies across similar problems.

Such analogical or associational reasoning is closely related to work on similarity-based recall, except that here the emphasis is not so much on which past experiences are more easily retrieved from memory but rather on which ones are more relevant. Relative to some of the literatures reviewed earlier, work on categorization and analogical reasoning is relatively underdeveloped in economics at this point.

Model-free learning. An extreme case of the idea that people simplify by “truncating” certain edges or nodes in the true data-generating process is to suppose that people entertain no causal mental representation of a problem at all, and purely learn from experience. This idea was popularized in the influential reinforcement learning framework in psychology that has found various applications in economics (e.g. Camerer and Hua Ho, 1999; Barberis and Jin, 2022). The main idea of these approaches is that people choose decisions that “worked in the past”, without having a structural understanding of why they worked.

Analogies: Learning from similar situations. A different strand of the literature, instead, captures the idea that people simplify through *problem substitution*: when people encounter a difficult problem that they cannot solve, they substitute it with a simpler problem that appears superficially similar. This line of work builds on much earlier research on attribute substitution (e.g., Kahneman et al., 2002). In this vein, Gilboa and Schmeidler (1995) proposed the idea of case-based reasoning, according to which people evaluate the utility of each potential action as the utility of the action in past decision environments, weighted by how similar these environments are to the present environment. As a result, the decision maker can be understood as taking decisions that “worked in the past”. This is a cognitive simplification strategy because it does not involve any explicit forward-looking cost-benefit calculations. Because case-based reasoning does not involve any explicit forward-looking mental representation of how actions map into payouts, this model is related in spirit to the variants of models of reinforcement learning discussed earlier (e.g., Camerer and Hua Ho, 1999; Barberis and Jin, 2022). The common theme of these models is that they do not involve an explicit forward-looking mental representation of how actions map into outcomes. The main difference is that accounts of reinforcement-learning do not rely on the logic of similarity and memory and are typically applied to relatively low-dimensional problems.

Cerigioni (2021) extends the Gilboa and Schmeidler (1995) model to allow for deliberation. In his model, when the current decision is sufficiently similar to one or more past decisions, the decision maker chooses what worked in the past. If, on the other hand, no similar situations come to mind, the decision maker rationally optimizes. This is framed as a dual decision process model. Webb et al. (2022) provide evidence for automatic decision-making in the field in the context of habit formation. Using a natural experiment (downsizing event), they document that a large share of persistence in consumption of canned tuna reflects automatic consumption, according to which people make purchasing decisions

“without actively thinking about it.”

In the model of policy uncertainty by Ilut and Valchev (2023) reviewed earlier, decision makers attempt to learn the policy function that maps a decision-relevant parameter into the optimal decision. In this model, decision maker learn from past observations, and such learning is similarity-based: they update more strongly about the local shape of the policy function at points in the parameter space that are more similar to those parameters that generated the empirical observation. This is a version of extrapolating from similar situations that is also modeled and studied empirically in Alsan et al. (2022). They argue that patients and doctors of color are more likely to learn from (or extrapolate from) clinical trials that included a larger number of patients of color.

Lumping together. In the class of analyses summarized above, people draw analogies between current and similar past situations. In a different class of models, people draw analogies between (and lump together) different current situations. Jehiel (2005) proposes a game-theoretic model in which people lump together different nodes at which other players must move, and then form expectations about behavior within each analogy class. Mullaithan et al. (2008) present a model of coarse thinking in the context of persuasion and Fryer and Jackson (2008) one of optimal categorization. An important feature of this entire class of models is that the categories are exogenously given and are not modeled to depend explicitly on a similarity function.

Evers et al. (2022) and Kőszegi and Matějka (2020) study mental accounting as a particularly important example of categorization. Evers et al. (2022) propose and document empirically that people are more likely to categorize different outcomes as belonging to the same “event” if the outcomes are more similar to each other. Kőszegi and Matějka (2020) endogenously derive mental accounting as an optimal response to information processing costs (rational inattention).

The experimental and empirical literature on categorization is still in its infancy. Thus far, most empirical work on memory has focused on settings in which memory shapes which input features enter people’s cognitive aggregation problems. An exception is Charles (2022b,a) who leverages the idea that companies may be categorized as “similar” in memory purely because they randomly appear next to each other on investor statements, or because they randomly announce earnings on the same day. As a result of these categorizations, when investors trade one stock, this increases the probability that the other stock comes to mind and is, hence, also traded. These papers document that this memory-based categorization has detectable impacts on market prices. I view the role of similarity in categorization as a promising area for empirical research.

7 Complexity Aversion

A final simplification strategy that the literature has discussed is systematic complexity aversion: people are said to systematically undervalue objects or choice sets that require a lot of information processing to evaluate. Loosely speaking, this can be understood as people being “risk averse” about the uncertainty that results from imperfect cognitive information processing (rather than from external uncertainty). Again, as in previous sections, this work can be partitioned into work on information-processing imperfections that result from cardinality and those that result from the difficulty of tradeoffs.

Cardinality. In choice under risk, the cardinality of a problem is governed both by the number of distinct payout states a lottery has and by the size of the choice set. There is a variety of experimental evidence to suggest that financial assets with many distinct payout states are systematically undervalued or avoided (Huck and Weizsäcker, 1999; Gillen et al., 2019; Carvalho and Silverman, 2019; Bernheim and Sprenger, 2020; Puri, 2022; Carvalho and Silverman, 2024). Iyengar and Kamenica (2010) document that a preference for simpler lotteries is especially pronounced when the choice set contains a larger number of options.

Fehr and Wu (2023) and references therein discuss evidence that people shy away from products that are overly complex due to obfuscation of add-on features. More generally, however, there is currently relatively little direct evidence on the importance of complexity aversion. One challenge here is to identify and test the different predictions that accounts of complexity aversion and accounts of behavioral attenuation make (because depending on the problem setup these can be similar).

On the theoretical side, the closest counterpart to this experimental evidence are models of complexity aversion (Puri, 2022; Hu, 2023). These typically have the form that a complexity cost function is added to the utility function, which depends on features such as support or entropy.

Tradeoffs. Models of caution often directly build on the idea – highlighted earlier – that tradeoffs are difficult (Cerreia-Vioglio et al., 2015, 2022; Chakraborty, 2021). In these models, the idea is that decision makers have subjective uncertainty about their utility from different choice options and act “cautiously” by assigning each option the minimum possible utility level, given the set of utility functions the decision maker deems possible. A less extreme version of this idea posits that people have uncertainty over their utils and take a concave transformation of utils, making them risk averse to uncertainty about their own utility weights. This literature argues that phenomena such as the endowment effect, the certainty effect and present bias reflect caution.

Viewed in combination, the literatures on behavioral attenuation and caution suggest that measures of noise and subjective uncertainty might be linked to behavior in two distinct

ways: (i) they might predict the *slope* of decision with respect to parameters (attenuation) and (ii) they might predict the *level* of decisions (caution). At this point, direct empirical evidence on these ideas would be very valuable, in particular on when which effect dominates. Moreover, while there is a growing body of both empirical and theoretical work on noisy information processing and resulting uncertainty and caution separately, I am not aware of models that capture both phenomena jointly.

8 What Determines Information Processing Imperfections?

An important presumption of much recent work is that there is no strong reason to expect behavioral anomalies to be constant across contexts. After all, if surprising choice behavior or beliefs largely reflect imperfect information processing, then we should see more pronounced anomalies when information processing costs are likely to be binding. This depends both on features of the problem and on characteristics of the decision maker.

8.1 The Complexity of Procedures and Algorithms

An obvious candidate mechanism that may partly drive the degree of imperfect information processing is complexity. Indeed, in computer science “complexity” refers to the cost of information processing. A main challenge is to be precise about what it is that makes problems more or less complex. Here, the literature can be partitioned according to whether it has attempted (i) to quantify the complexity of a problem (or choice option) based on objective features of the problem; or (ii) to quantify the complexity of the procedures and algorithms people may actually use to solve a problem. Naturally, some of the approaches pursued in the literature can be interpreted to fall into either of these buckets.

In Sections 2 and 3.2, I already discussed which features make decision problems complex, focusing in particular on complexity-from-dissimilarity (the strength of tradeoffs) and cardinality. I now discuss work on what makes procedures complex.

A broad idea in recent work on complexity is that when decision problems are reasonably simple, behavior can be well-approximated by the standard maximizing model but that when problems are more complex, people instead rely on specific procedures or algorithms to solve problems that may be intractable otherwise (Camara, 2022). According to this idea, people may entertain a “library” of specific algorithms and problem-solving approaches and deploy one of them for any given problem, rather than attempt to maximize.¹⁰

Evidence on algorithmic decision making. In support of the general idea that complexity causes people to deploy more algorithmic decision processes, Arrieta and Nielsen (2023)

¹⁰There is a long tradition in economics of directly modeling certain decision procedures and algorithms. Much of the earlier literature on specific algorithms relies on models of automata (Rubinstein, 1998).

document that as lottery choice problems become more complex, people perform better in describing their decision making process to others who need to replicate their decisions. This is interpreted as evidence that complexity causes the usage of specific, simple procedures.

Related recent experimental work has begun to shed light on which algorithms are more or less complex to implement. Oprea (2020) and Banovetz and Oprea (2022) experimentally test some of the ideas in the theoretical automata literature. For example, Oprea (2020) elicits people's willingness-to-pay to avoid having to implement certain rules as a measure of how complex different algorithms are. A main takeaway from these papers is that people avoid implementing rules that have a higher "dimensionality" (e.g., a larger number of states or transitions), consistent with the discussion of dimensionality complexity above.

Metrics of computational complexity. An important question is which algorithmic rules are too complex to be implemented, and what corresponding complexity metrics are. This work is partly influenced by complexity notions in computer science. All of the recent papers I am aware of formalize notions of complexity that – loosely speaking – partly reflect the dimensionality of the problem.

Camara (2022) shows that the very weak assumption that people cannot solve NP-hard problems implies that they have to engage in narrow bracketing (also see Bossaerts and Murawski, 2017). Sanjurjo (2023) models the complexity of algorithmic rules as arising from space complexity (as in computer science), and tests this idea using multi-attribute choice data. Both of these models predict that larger choice sets are more complex.

Salant and Spenkuch (2022) model an algorithmic satisficing procedure according to which the decision maker randomly samples a subset of alternatives (the consideration set) and then generates a noisy evaluation of them, where the noisiness of the evaluation increases in complexity. Here, the complexity of an option is modeled as the depth of the sub-game that it induces. Salant and Spenkuch (2022) provide empirical support for such a model using data on chess.

When are comparisons difficult? Dissimilarity and the strength of tradeoffs. As discussed in Section 4, comparative thinking is a main response to complexity. Yet when are comparisons easy to make? In Section 3, I discussed work that suggests that *dimension-by-dimension dissimilarity causes complexity* because it makes comparisons difficult.

The work on dissimilarity causing complexity can be interpreted in two different ways. First, it formalizes and quantifies the complexity of a choice problem. Second, it formalizes the complexity of a specific cognitive algorithm (that of comparing dimension-by-dimension and then aggregating the comparisons).

Takeaways. Overall, research on what makes decision problems complex is still in its infancy. Economists have devoted substantially more effort to quantifying the magnitude (and

costs) of inattention than to quantifying the complexity of a problem. While this is beginning to change, most of the work on what makes a problem complex falls either in the dimensionality or dissimilarity complexity bucket. Similarly, we are only beginning to understand which specific algorithms people may deploy when they are confronted with complex problems, and how we can quantify the complexity of these algorithms themselves.

8.2 Features of the Decision Maker

Availability of cognitive resources. There is much evidence to suggest that the availability of cognitive resources has substantial effects on observed choices and beliefs. This is difficult to reconcile with a perspective of choice anomalies reflecting preferences but entirely natural from a perspective of imperfect information processing.

For example, various studies document that time pressure or cognitive load interventions have strong effects on intertemporal decision making (Ebert and Prelec, 2007; Benjamin et al., 2013; Imas et al., 2021). Similarly, a large literature documents effects of cognitive load on choice under risk (e.g., Gerhardt et al., 2016). The results from these experimental interventions are consistent with a large number of correlational studies that have linked measures of cognitive ability to choice behavior in domains such as impatience (Benjamin et al., 2013; Dohmen et al., 2010; Falk et al., 2018), risk aversion (Dohmen et al., 2010) and probability weighting (Choi et al., 2022).

Experience. Multiple recent papers document that behavioral attenuation weakens as decision makers gather more experience with a specific problem configuration. These results have been documented in contexts such as risk aversion (Frydman and Jin, 2022), probability weighting (Frydman and Jin, 2023) and coordination games (Frydman and Nunnari, 2023). This is typically interpreted through the lens of models of efficient coding – the idea that the brain optimally reduces noise (and hence attenuation) for problems that it encounters more frequently.

9 Revisiting Some Classics

In my experience, at a deep level, behavioral economists always knew that many of their reduced-form models of non-standard utility functions and parametric biases weren't literally true. Too much evidence accumulated over time that suggested that classic behavioral phenomena vary strongly across contexts and elicitation methods, and with the availability of cognitive resources. Moreover, the pioneers of the field were sometimes fairly explicit in their writings that their explanations and models were meant to be reduced-form adaptations of ideas in cognitive psychology (e.g., Kahneman and Tversky, 1979). Nonetheless, once reduced-form models are proposed and widely adopted, they are sometimes taken to

be true in a more literal sense than the pioneers intended. As I summarize now, the cognitive turn has begun to make meaningful progress on reducing quite a few of behavioral economists' greatest success stories to a smaller set of cognitive primitives related to imperfect information processing. Much work remains to be done, but I believe the path is now clearer.

In discussing the origins of behavioral economics classics, I largely stick to the empirical evidence. In other words, I do not declare success for the cognitive turn in explaining a certain phenomenon purely because a cognitive model explains the pattern, but only when direct empirical evidence supports such an interpretation.

9.1 What We (Probably) Understand

Probability weighting. Probability weighting is a special case of behavioral attenuation. In my opinion, the evidence that links it to noisy information processing (Enke and Graeber, 2023; Frydman and Jin, 2023; Oprea, 2022; Vieider, 2022; Enke et al., 2024a) or cognitive resources (Choi et al., 2022), and that explains how complexity (dissimilarity) and the decision format determine when it does (not) appear (Shubatt and Yang, 2024) is numerous and internally consistent.

Choice-versus-valuation preference reversals. A long literature documents preference reversals between choice and valuation, such as when people choose a relatively safe gamble over a long-shot one, but when asked to value the two independently they value the long-shot one higher. These patterns largely reflect the difficulty of navigating tradeoffs. Intuitively, valuing a long-shot gamble (such as “getting \$120 with probability 10%”) in terms of a certainty equivalent is cognitively very difficult because the “probability currency” of 10% is very different from that of 100%. In contrast, choosing between the two lotteries is easier because the payout probabilities are more similar to each other (Shubatt and Yang, 2024).

Hyperbolic discounting (not dynamic inconsistency). There are two classes of hyperbolic discounting in the literature. A first is time inconsistency and failures of self-control, chiefly over primary rewards. Thus far, there is little or no evidence in the cognitive foundations literature that suggests that lack of self-control and dynamic inconsistency is related to imperfect information processing. A second class is the hyperbolicity of the discount function over long horizons, chiefly over financial flows. This second pattern is a special case of behavioral attenuation. Multiple papers have documented that hyperbolicity strongly increases with proxies for noisy information processing (Enke et al., 2023a, 2024a), time pressure (Ebert, 2001; Ebert and Prelec, 2007), and that it can be predictably switched on and off as a function of option dissimilarity and complexity (Shubatt and Yang, 2024).

Scope insensitivity. A huge literature documents scope insensitivity in contingent valuation studies (Diamond and Hausman, 1994). For quite some time, researchers interpreted these patterns as potentially reflecting heavily non-linear preferences – maybe people actually value saving 100 birds roughly as much as saving 10,000 birds? Recent work has shown that scope insensitivity is a special case of behavioral attenuation – it is strongly correlated with cognitive uncertainty and decreases in experimental interventions that make evaluations easier (Toma and Bell, 2022; Enke et al., 2024a).

Decoy and compromise effects. The literature contains a large number of demonstrations of a “preference for compromise”, in studies involving decoys and others. While the jury is still out on what exactly the correct explanation for these patterns is, it is now clear that the answer is neither preferences nor fixed heuristics. Rather, a growing body of evidence highlights that these effects are also driven by information processing imperfections, in particular comparative thinking (Bushong et al., 2021; Somerville, 2022) and comparison complexity resulting from dimension-by-dimension dissimilarity (Natenzon, 2019; Shubatt and Yang, 2024).

Over- and underreaction. Over- and underreaction in belief updating are classic behavioral economics phenomena but until recently it wasn’t understood when which of the two should prevail. Research on cognitive foundations has again emphasized that understanding these phenomena from a perspective of information processing is productive. We now understand that underreaction is largely a special case of behavioral attenuation (Enke and Graeber, 2023; Augenblick et al., 2021; Ba et al., 2022), while overreaction can be driven by a combination of behavioral attenuation and incorrect problem representations resulting from limited attention and memory (Augenblick et al., 2021; Ba et al., 2022; Enke et al., 2024b; Bordalo et al., 2023a).

Compressed forecasts and expectations. A large number of papers in behavioral economics and the social economics survey literature document that beliefs, perceptions, expectations and forecasts are typically strongly compressed functions of the underlying truth, giving rise to patterns such as “everyone is middle class”. There is now a bulk of evidence that suggests that these patterns are another special case of noisy information processing and resulting behavioral attenuation (Fischhoff and Bruine De Bruin, 1999; Enke and Graeber, 2023; Giustinelli et al., 2019).

Full neglect and multi-modality. While many behavioral anomalies can be characterized through a “partial neglect” or “anchoring-and-adjustment” logic, some others reflect the full neglect of certain problem aspects, as evidenced by strongly multi-modal response patterns in work on belief updating. These reflect incorrect problem representations that are usually

driven by selective attention to some problem features (Enke and Zimmermann, 2019; Enke, 2020; Graeber, 2022; Fan et al., 2023; Bordalo et al., 2023a). This is true for phenomena such as base rate neglect, correlation neglect, selection neglect, and others.

Stereotyping: exaggerated differences. Stereotyping reflects a special case of comparative thinking, according to which people directly compare the relative plausibility of different hypotheses and frequencies, which makes them overreact relative to the true differences that exist (Coffman, 2014; Bordalo et al., 2016, 2019, 2021).

Description-experience gap. The widely-studied difference between risk-taking based on description (direct presentation of probabilities) and experience (sequentially sampling from the outcome distribution) is largely driven by salience and memory (Bohren et al., 2024). In decision-making based on description, unlikely but salient large outcomes attract attention, while in decisions from experience rare extreme outcomes tend to be forgotten, producing behavior that looks like more pronounced probability weighting in decisions based on description.

9.2 Where More Work is Needed

Endowment effect. One view of the endowment effect is that it reflects loss aversion. Another view suggests that it is driven by the difficulty of translating product attributes into a common currency (such as dollars or utils), see Cerreia-Vioglio et al. (2022). Much indirect evidence suggests that the endowment effect has cognitive origins – it decreases with experience, disappears when people know the market value of goods, and is uncorrelated with experimental measures of loss aversion (Chapman et al., 2023b). It would be very useful to directly test the idea that the endowment effect emerges because people have a hard time valuing objects (and potentially exhibit caution / a type of risk aversion over their own valuation uncertainty).

Reference points. There can be no question that people often compare outcomes to reference points, and much of the discussion in behavioral economics has focused on how exactly such reference points should be modeled. But at a fundamental level, we still don't even know whether comparing outcomes with a reference points reflects true reference-dependent preferences or is a result of information-processing imperfections (O'Donoghue and Sprenger, 2018). For example, in light of the emphasis on the difficulty of across-dimension aggregation in recent work, comparing outcomes to a reference point within dimensions may be easier than to aggregate across dimensions.

Allais paradoxes. There is a long discussion about whether the Allais paradoxes represent information-processing mistakes or preferences (e.g., Rubinstein, 1988; McGranaghan et al.,

2022; Cerreia-Vioglio et al., 2015; Bordalo et al., 2012). From an empirical perspective, I believe we still don't know the answer.

Present bias. As noted above, behavioral attenuation is most likely a driver of the hyperbolic shape of the empirical discount function, but not of present bias or dynamic inconsistency. At the same time, various models (Rubinstein, 2003; Fudenberg and Levine, 2006; Chakraborty, 2021) and some evidence (Imas et al., 2021; Chakraborty et al., 2017) suggest that cognitive mechanisms such as self-control, valuation uncertainty about the future, similarity or noise can contribute to observed present bias.

Utility curvature. Time and again, both lab and field data suggest implausibly high levels of estimated utility curvature (Rabin, 2000). While behavioral economists routinely estimate non-linear utility functions, a large majority appears to agree that these estimates do not really reflect true diminishing marginal utility. As discussed in Section 3, one possibility is that seeming diminishing marginal utility is a special case of diminishing sensitivity as generated by noisy information processing. Intuitively, people understand that \$80 and \$5 are both better than \$0, yet how much more is cognitively difficult to assess, and might lead to a compression effect that looks like diminishing marginal utility (see the discussion in Kahneman and Tversky, 1979). It would be very helpful if more direct evidence on this could be gathered. This said, I do not wish to suggest that *all* of utility curvature isn't real – some almost certainly is, but some I suspect isn't.

Context effects, analogical and similarity-based reasoning. As discussed above, a plausible – and I believe in reality widespread – simplification strategy is to make decisions based on what worked well in similar decision problems in the past. Here, the emphasis is on “similar problems” – it is, of course, widely understood based on the reinforcement learning literature that people make decisions based on what worked well in exactly the same problem in the past. However, I believe that much is to be gained from empirical work that attempts to understand “spillovers” and mental analogies across problems (Gilboa and Schmeidler, 1995; Bordalo et al., 2023b).

9.3 What Seems Largely Unrelated

Social preferences and fairness. There is little evidence to suggest that mistakes have much to do with the existence of social motivations such as those related to prosociality, fairness, cooperation, moral wiggle room behavior and motivated reasoning. This being said, while information-processing imperfections do not cause the *existence* of social motivations, a considerable body of evidence suggests that they affect how *social motivations translate into behavior*. For example, public goods giving is significantly affected by confusion and resulting behavioral attenuation (e.g., Andreoni, 1995; Bao and Pei, 2023). Similarly, I speculate

that the widespread equal split fairness norm in parts reflect a simplification strategy – when people find it difficult to assess who contributed how much to a common project, splitting the pie equally is a simple thing to do.

Belief-based utility. Much work suggests that people have preferences over their beliefs (see, e.g., Loewenstein and Wojtowicz, 2023, for a review). This includes work on news utility, anticipation, ego, social image, self image, moral wiggling, curiosity, and others. Again, as I see it, there is little reason to believe that these have cognitive origins. However, as is the case with social preferences, information processing limitations can act as an “enabler” for moral wiggle room behavior or motivated reasoning because it makes the mapping between intrinsic motivations and behaviors less transparent (e.g. Haisley and Weber, 2010; Exley and Kessler, 2019).

10 Evidence From the Field

The insights derived from the cognitive turn have started to become visible in applied economics. I here showcase a few examples to illustrate the breadth of applications, rather than comprehensively summarize the applied literature. The applications summarized below span papers in (macro-) finance, labor and development economics.

Cognitive noise and behavioral attenuation. Augenblick et al. (2021) study the role of behavioral attenuation for market-implied beliefs as reflected in option prices for the S&P500. They study whether traders (and resulting prices) are attenuated to the precision of the information in the market. When an option is far from expiration, the information in the market is weak and when it is close to expiration, the signals are very informative. Attenuation predicts that option prices move too much when the information is weak (far from expiration) yet too little when the information is strong (close to expiration). Augenblick et al. (2021) document that this is indeed the case, reconciling the joint existence of over- and underreaction of expectations, purely as a mechanical outcome of attenuation.

Drerup et al. (2017) and Giglio et al. (2021) study the so-called attenuation puzzle in stock market participation: the puzzle that the link between future return expectations and equity share is quantitatively much too small relative to canonical models, even after classical measurement error in beliefs is instrumented out. Both of these papers present evidence suggesting that this attenuation puzzle is a special case of noise-driven behavioral attenuation. Both papers combine administrative data on stock holdings with survey data on expectations. Crucially, both papers also measure people’s cognitive uncertainty (or confidence in) the meaningfulness of their stated belief distributions. The results show that the attenuation puzzle is strongly concentrated in people who exhibit high uncertainty about their beliefs.

Relatedly, D'Acunto et al. (2023) study attenuation effects in consumers' response to the incentives resulting from changes in government programs. They document that the field decisions of higher cognitive ability individuals are more responsive to variation in incentives, which affects behavior in contexts ranging from borrowing to car purchases.

Card et al. (2024) study the role of cognitive noise in the decision-making of economics journal editors. Just like the many other contexts emphasized in this paper, editors' decisions a potentially difficult process of mentally aggregating and trading off different signals (such as conflicting referee reports). The authors leverage the key insight from sequential evidence accumulation models that longer response times are associated with closer proximity to indifference, and use data on editorial decision times to predict outcomes. Consistent with standard sequential evidence accumulation models, they find that conditional on an R&R, editor response times are negatively predictive of paper quality, while conditional on a Reject decision, response times are positively predictive of paper quality.

Complexity. Abeler et al. (2023) document the role of complexity in workers' incentive contracts. They consider a setting in which a firm announces that it will use current productivity as a target for future bonus payments, a classic example of a setup that should give rise to a Ratchet effect (workers reducing current effort to decrease the future target). Yet the authors show that because the incentive scheme is described in a sufficiently complex way, the Ratchet effect is almost entirely absent. In contrast, when the incentives are laid out in simpler terms, workers' current effort exhibits the canonical Ratchet response.

Goal-directed attention. Bartoš et al. (2016) deploy tools to measure information acquisition in a field experiment and document that asymmetric information search as predicted by rational inattention contributes to discrimination against minority groups. The main idea is that when a decision maker entertains the strong prior belief that members of a certain group are of low "quality", on average, optimal information search will downweight or even ignore this group if the decision maker is only interested in identifying top candidates. Conversely, if the decision maker is only interested in avoiding the worst candidates, information search will focus on the minority group. The authors test this idea in two field experiments, by tracking the online information acquisition of employers and landlords who evaluate applications. The authors show that information acquisition is in line with the predictions of directed attention models and contributes to systematic discrimination of minorities.

Memory. The by-now classic evidence on experience effects (Malmendier and Nagel, 2011, 2016) was some of the first evidence that directly pointed at the importance of memory for the formation of subjective expectations in economics. In recent years, various papers have studied the role of associative memory in finance. Charles (2022a) documents that contextual associations affect trading behavior. The idea is that when two stocks are associated

in memory (because they appear next to each other on portfolio statements due to their alphabetical ranking), the cueing of one stock induces the investor to also trade the other one. Charles (2022b) shows that memory-induced trading patterns affect market prices. The paper leverages a similar idea, which is that companies may be associated in memory when they have overlapping earnings announcement schedules. The paper shows that when the announcements of two firms overlap, the separate earnings announcement of one of the firms several months later produces buying pressure in the other firm's stock. The interpretation is that because the two firms are associated in memory, the announcement of one firm directs attention also to the other one even when it does not announce earnings.

Jiang et al. (2022) document the importance of associative recall by combining surveys of investors' memories with administrative trading data. The headline result is that when market returns are positive, people become more likely to remember episodes of rising markets from the past, a hallmark signature of associative recall. These investor memories of past returns, in turn, are strongly predictive of future return expectations and extrapolation.

Augenblick et al. (2023) study the importance of imperfect and associative recall in a development context, by studying the savings patterns of Zambian farmers. They document that people have a pronounced tendency to forget about upcoming expenditures that are small, irregular and stochastic, and as a result save too little. To improve recall, they leverage the idea of associative recall by asking farmers to think through their expenditures through categories. This intervention strongly cues the retrieval of associated expenditure events, and has large effects on savings behavior.

11 Outlook and Open Questions

There is a large number of open questions in the cognitive turn, some theoretical, some experimental and some empirical in nature.

11.1 Theory

While the cognitive turn has arguably made some progress on reducing the large number of distinct behavioral anomalies into a shorter list of simplification strategies, it is equally true that at this point there is no single "plug-and-play" model of simplification strategies that can be immediately injected into field applications. This is partly a theoretical and partly an empirical challenge. On the theoretical side, the different literatures that I have summarized in this review – such as those on noisy cognition and stimulus-driven attention or recall – have largely evolved in isolation from each other. A first-order challenge is to articulate how they are interrelated, when which one plays out, and whether a single model can capture all (or most) of them.

Moreover, even within each of these literatures important challenges remain. For ex-

ample, there is no consensus on which exact assumptions researchers should place on the location of the noise (or the prior) in Bayesian cognitive noise models. Similarly, there is no consensus on the question of which features stimulus-driven attention and memory are supposed to operate on. These are important challenges because if a different type of noise or a different cue need to be invoked for each anomaly, then behavioralists are simply transforming a proliferation of reduced-form concepts into a proliferation of cognitive foundations. Closely related, on the empirical side further work is needed to make cognitive foundations models broadly applicable because – unlike canonical models of reference-dependence, present bias or social preferences – they are not only defined over payoffs but also over primitives such as complexity or contextual similarity that are traditionally not measured by economists.

11.2 Economic Applications

Multi-agent settings. Up to this point, research on the cognitive turn is largely restricted to individual decision-making. For example, the emergence of the cognitive turn contributed to transforming experimental economics away from experiments on markets and strategic interactions towards individual decision experiments. Yet sooner rather than later, behavioralists and experimentalists alike will have to show that the insights they've developed in individual decision contexts also illuminate more traditional objects of economic interest. In this regard, a key challenge will be to move to multi-agent settings without giving up the insistence on correct cognitive foundations (for some examples along these lines see Frydman and Nunnari, 2023; Enke et al., 2023b; Amelio, 2023; Li and Camerer, 2022).

More field applications. A large majority of the papers summarized in this review are theoretical contributions, lab and online experiments, or surveys. Future work on the topic will increasingly document the relevance of cognitive foundations for ecological behaviors. A key challenge in this regard is measurement: as noted above, unlike canonical traditional behavioral economics models, cognitive foundations models are often not only defined over payoffs but also over primitives such as complexity, noise or contextual similarity that are traditionally not measured by economists. To overcome this problem, researchers have (and I predict increasingly will) rely on a combination of non-choice data and machine learning techniques.

The cognitive turn has used non-choice data to gather evidence on cognitive foundations, and to link them to behaviors (see the discussion by Camerer, 2008). This includes elicitation of descriptions of procedures (Arrieta and Nielsen, 2023) or of choices between axioms (Nielsen and Rehbeck, 2022), the measurement of cognitive uncertainty (Enke and Graeber, 2023), the constructive use of response time data (Rubinstein, 2007, 2016; Alós-Ferrer et al., 2021; Liu and Netzer, 2023; Card et al., 2024) and the measurement of information acqui-

sition (e.g., Bartoš et al., 2016; Caplin and Dean, 2015). Techniques like eye-tracking might well re-gain popularity in the wake of this movement (e.g., Arieli et al., 2011; Reutskaja et al., 2011; Engelmann et al., 2021).

If experimentalists increasingly opt for more naturalistic designs, then machine learning techniques will become more important in quantifying the relevance of intrinsically qualitative features such as text, speech or visual features. For example, Graeber et al. (2023a,b) show how speech data can be incorporated into experiments to study information transmission and the spread of narratives. Li and Camerer (2022) and Bose et al. (2022) provide evidence that a standard machine learning algorithm from neuroscience that predicts the visual salience of portions of images helps to predict behavior in standard consumer choice and stock market forecasting contexts.

11.3 Experiments

The cognitive turn was partly driven by new experimental paradigms, but it also suggests productive avenues for further change going forward. Traditionally, experimental economists designed decision environments that were very abstract. The rationale was to strip away all context that could potentially “pollute” the identification of the theoretical construct that the researcher sought to measure. Experiments were only considered “clean” if they were stripped down to the bare minimum that the corresponding economic model posited, but were fully specified as far as economic incentives and procedures are concerned. Elicitation mechanisms for beliefs or preferences were considered appropriate if and only if they satisfied theoretical incentive-compatibility properties.

While minimalism, incentive compatibility, objectivity and unconfoundedness are no doubt desirable properties for experiments to have, it is becoming increasingly clear that the way in which these goals were achieved implied non-trivial costs, and itself created other confounds, as I discuss in the remaining paragraphs.

Simplification strategies confound measurements of preferences and beliefs. Perhaps the most important such cost is that abstract and unfamiliar study designs may not elicit the true, “unpolluted” construct the researcher seeks to measure (or that an economic model specifies). Instead, such designs may produce behavior that reflects confusion and heuristics. Importantly, given the often unnatural and abstract nature of experimental economics designs, there is no strong a priori reason to expect that those simplification strategies that participants resort to in experiments are also the ones they deploy in environments that are more naturalistic or with which they have more experience. In this spirit, a large number of contributions have documented that commonly-used experimental paradigms can produce severely biased measures of preferences and beliefs (e.g. Enke et al., 2023a; Andersson et al., 2020; Bouchouicha et al., 2023), and we are only beginning to understand how these short-

comings can be improved (e.g., Alós-Ferrer et al., 2021; Kendall and Chakraborty, 2022; Halevy et al., 2023).

An important special case of this class of problems is that of scoring rules. Traditionally, theorists and experimentalists favored scoring rules that are incentive compatible under the assumption of perfect rationality. However, recent work has shown that people systematically misunderstand complex scoring rules and “hedge” against their confusion (Danz et al., 2022). As a result, experimentalists’ emphasis on incentive compatibility can sometimes produce internally contradictory research designs. For example, suppose a researcher tests the hypothesis that a subject makes systematic errors in processing risky lotteries – under this hypothesis, it is not clear why a subject should understand that a scoring rule – that often constitutes a significantly more complicated lottery than the object the subject is being asked to value – is incentive compatible. Similarly, economists insist on deploying ingenious scoring rules to elicit beliefs and expectations, even when they test the hypothesis that people fall prey to biases such as base rate neglect – yet it is arguably not at all obvious that rationally incorporating a base rate is more difficult than thinking through the incentive compatibility properties of highly non-trivial scoring rules. A challenge going forward will be to devise scoring rules and other experimental procedures that are incentive compatible also under information processing constraints.

Contextual cues and naturalism: Tossing out the baby with the bathwater. A second cost that is entailed by making experimental paradigms minimalistic and abstract is that – essentially by assumption – it negates a potential role for the contextual cues that many researchers today believe are important for understanding the context-dependence of behavior. This suggests that future research will benefit from bringing back (in a controlled form) some of those contextual elements that standard elicitation protocols have sought to eliminate.

Another example is that economists’ desire to make choice environments well-defined has led them to almost exclusively rely on experimental environments in which the choice set is fully and transparently displayed to the decision maker (“Do you prefer this pizza or the pasta?”). Yet in reality people need to construct their consideration sets from memory (“What could I have for dinner? There’s the Italian place, the Vietnamese one...”), which naturally gives rise to a role for memory and cues. I believe that much is to be gained from developing more naturalistic choice environments without giving up too much control.

References

- ABELER, J., A. FALK, L. GOETTE, AND D. HUFFMAN (2011): “Reference points and effort provision,” *American Economic Review*, 101, 470–492.
- ABELER, J., D. B. HUFFMANN, AND C. RAYMOND (2023): “Incentive Complexity, Bounded Rationality and Effort Provision,” .
- ABELER, J. AND S. JÄGER (2015): “Complex Tax Incentives,” *American Economic Journal: Economic Policy*, 7, 1–28.
- AFROUZI, H., S. Y. KWON, A. LANDIER, Y. MA, AND D. THESMAR (2023): “Overreaction in expectations: Evidence and theory,” *The Quarterly Journal of Economics*, qjad009.
- AGRANOV, M., P. J. HEALY, AND K. NIELSEN (2020): “Non-Random Randomization,” *Available at SSRN*.
- AGRANOV, M. AND P. ORTOLEVA (2017): “Stochastic choice and preferences for randomization,” *Journal of Political Economy*, 125, 40–68.
- (2020): “Ranges of Preferences and Randomization,” *Working Paper*.
- AGRANOV, M. AND P. RESHIDI (2023): “Disentangling Suboptimal dating: Complexity, Structure, and Sequencing,” .
- AINA, C. (2021): “Tailored Stories,” Tech. rep., Mimeo.
- ALÓS-FERRER, C., E. FEHR, AND N. NETZER (2021): “Time will tell: Recovering preferences when choices are noisy,” *Journal of Political Economy*, 129, 1828–1877.
- ALSAN, M., M. DURVASULA, H. GUPTA, J. SCHWARTZSTEIN, AND H. L. WILLIAMS (2022): “Representation and Extrapolation: Evidence from Clinical Trials,” Tech. rep., National Bureau of Economic Research.
- AMBUEHL, S. AND S. LI (2018): “Belief updating and the demand for information,” *Games and Economic Behavior*, 109, 21–39.
- AMELIO, A. (2023): “Social Learning, Behavioral Biases and Group Outcomes,” .
- AMERIKS, J., G. KÉZDI, M. LEE, AND M. D. SHAPIRO (2020): “Heterogeneity in expectations, risk tolerance, and household stock shares: The attenuation puzzle,” *Journal of Business & Economic Statistics*, 38, 633–646.
- ANDERSSON, O., H. J. HOLM, J.-R. TYRAN, AND E. WENGSTRÖM (2016): “Risk aversion relates to cognitive ability: Preferences or noise?” *Journal of the European Economic Association*, 14, 1129–1154.

- (2020): “Robust inference in risk elicitation tasks,” *Journal of Risk and Uncertainty*, 61, 195–209.
- ANDRE, P. (2022): “Shallow meritocracy,” *Available at SSRN 3916303*.
- ANDRE, P., I. HAALAND, C. ROTH, AND J. WOHLFART (2023a): “Narratives about the Macroeconomy,” .
- ANDRE, P., C. PIZZINELLI, C. ROTH, AND J. WOHLFART (2022): “Subjective models of the macroeconomy: Evidence from experts and representative samples,” *The Review of Economic Studies*, 89, 2958–2991.
- ANDRE, P., P. SCHIRMER, AND J. WOHLFART (2023b): “Mental models of the stock market,” .
- ANDREONI, J. (1995): “Cooperation in public-goods experiments: kindness or confusion?” *The American Economic Review*, 891–904.
- ANDREONI, J. AND C. SPRENGER (2011): “Uncertainty equivalents: Testing the limits of the independence axiom,” Tech. rep., National Bureau of Economic Research.
- ARIELI, A., Y. BEN-AMI, AND A. RUBINSTEIN (2011): “Tracking decision makers under uncertainty,” *American Economic Journal: Microeconomics*, 3, 68–76.
- ARIELY, D., G. LOEWENSTEIN, AND D. PRELEC (2003): ““Coherent arbitrariness”: Stable demand curves without stable preferences,” *The Quarterly journal of economics*, 118, 73–106.
- ARRIETA, G. AND K. NIELSEN (2023): “Procedural Decision-Making in the Face of Complexity,” *Working Paper*.
- AUGENBLICK, N., B. K. JACK, S. KAUR, F. MASIYE, AND N. SWANSON (2023): “Retrieval Failures and Consumption Smoothing: A Field Experiment on Seasonal Poverty,” .
- AUGENBLICK, N., E. LAZARUS, AND M. THALER (2021): “Overinference from Weak Signals and Underinference from Strong Signals,” *arXiv preprint arXiv:2109.09871*.
- AZEREDO DA SILVEIRA, R. AND M. WOODFORD (2019): “Noisy memory and over-reaction to news,” in *AEA Papers and Proceedings*, vol. 109, 557–61.
- BA, C., J. A. BOHREN, AND A. IMAS (2022): “Over-and Underreaction to Information,” *Available at SSRN*.
- BADDELEY, A. (2013): *Essentials of human memory (classic edition)*, Psychology Press.

- BAKHTIN, M., M. NIEDERLE, AND M. MÖBIUS (2023): “A Memory Model of Belief Formation,” *Working Paper*.
- BANOVETZ, J. AND R. OPREA (2022): “Complexity and Procedural Choice,” *American Economic Journal: Microeconomics*.
- BAO, T. AND J. PEI (2023): “Cognitive Uncertainty, GPT, and Contribution in Public Goods Game,” *Working Paper*.
- BARBERIS, N. AND L. JIN (2022): “Model-free and Model-based Learning as Joint Drivers of Investor Behavior,” Tech. rep., Mimeo.
- BARLOW, H. B. (1961): “Possible principles underlying the transformation of sensory messages,” *Sensory communication*, 1, 217–233.
- BARRON, K. AND T. FRIES (2023): “Narrative persuasion,” .
- BARRON, K., S. HUCK, AND P. JEHIEL (2023): “Everyday econometricians: Selection neglect and overoptimism when learning from others,” .
- BARTOŠ, V., M. BAUER, J. CHYTILOVÁ, AND F. MATĚJKA (2016): “Attention discrimination: Theory and field experiments with monitoring information acquisition,” *American Economic Review*, 106, 1437–1475.
- BAYRAK, O. K. AND J. D. HEY (2020): “Understanding preference imprecision,” *Journal of Economic Surveys*, 34, 154–174.
- BEAUCHAMP, J. P., D. I. LAIBSON, AND C. F. CHABRIS (2019): “Measuring and controlling for the compromise effect when estimating risk preference parameters,” *Experimental Economics*, 1–31.
- BELZIL, C. AND T. JAGELKA (2020): “Separating True Preferences from Noise and Endogenous Effort,” .
- BENJAMIN, D. J. (2019): “Errors in Probabilistic Reasoning and Judgmental Biases,” in *Handbook of Behavioral Economics*.
- BENJAMIN, D. J., S. A. BROWN, AND J. M. SHAPIRO (2013): “Who is ‘behavioral’? Cognitive ability and anomalous preferences,” *Journal of the European Economic Association*, 11, 1231–1255.
- BERNHEIM, B. D. AND A. RANGEL (2004): “Addiction and cue-triggered decision processes,” *American economic review*, 94, 1558–1590.

- BERNHEIM, B. D. AND C. SPRENGER (2020): “On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting,” *Econometrica*, 88, 1363–1409.
- BHARGAVA, S. AND R. FISMAN (2014): “Contrast effects in sequential decisions: Evidence from speed dating,” *Review of Economics and Statistics*, 96, 444–457.
- BHATIA, S. AND G. LOOMES (2017): “Noisy preferences in risky choice: A cautionary note.” *Psychological review*, 124, 678.
- BODOH-CREED, A. L. (2020): “Mood, memory, and the evaluation of asset prices,” *Review of Finance*, 24, 227–262.
- BOHREN, A., J. HASCHER, A. IMAS, M. UNGEHEUER, AND M. WEBER (2024): “A Cognitive Foundation for Perceiving Uncertainty,” *Available at SSRN*.
- BOHREN, J. A. AND D. N. HAUSER (2021): “Learning with heterogeneous misspecified models: Characterization and robustness,” *Econometrica*, 89, 3025–3077.
- BORDALO, P., G. BURRO, K. B. COFFMAN, N. GENNAIOLI, AND A. SHLEIFER (2022a): “Imagining the future: Memory, simulation and beliefs about covid,” Tech. rep., National Bureau of Economic Research.
- BORDALO, P., K. COFFMAN, N. GENNAIOLI, F. SCHWERTER, AND A. SHLEIFER (2021): “Memory and representativeness.” *Psychological Review*, 128, 71.
- BORDALO, P., K. COFFMAN, N. GENNAIOLI, AND A. SHLEIFER (2016): “Stereotypes,” *Quarterly Journal of Economics*, 131, 1753–1794.
- (2019): “Beliefs about gender,” *American Economic Review*, 109, 739–73.
- BORDALO, P., J. CONLON, N. GENNAIOLI, S. KWON, AND A. SHLEIFER (2023a): “How People Use Statistics,” *Working Paper*.
- BORDALO, P., J. J. CONLON, N. GENNAIOLI, S. Y. KWON, AND A. SHLEIFER (2023b): “Memory and probability,” *The Quarterly Journal of Economics*, 138, 265–311.
- BORDALO, P., N. GENNAIOLI, Y. MA, AND A. SHLEIFER (2020a): “Overreaction in macroeconomic expectations,” *American Economic Review*, 110, 2748–82.
- BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2012): “Salience theory of choice under risk,” *Quarterly Journal of Economics*, 127, 1243–1285.
- (2013): “Salience and Consumer Choice,” *Journal of Political Economy*, 121, 803–843.

- (2018): “Diagnostic expectations and credit cycles,” *The Journal of Finance*, 73, 199–227.
- (2020b): “Memory, attention, and choice,” *The Quarterly journal of economics*.
- (2022b): “Overreaction and diagnostic expectations in macroeconomics,” *Journal of Economic Perspectives*, 36, 223–244.
- (2022c): “Saliency,” *Annual Review of Economics*, 14, 521–544.
- BOSE, D., H. CORDES, S. NOLTE, J. C. SCHNEIDER, AND C. F. CAMERER (2022): “Decision weights for experimental asset prices based on visual saliency,” *The Review of Financial Studies*, 35, 5094–5126.
- BOSSAERTS, P. AND C. MURAWSKI (2017): “Computational complexity and human decision-making,” *Trends in cognitive sciences*, 21, 917–929.
- BOUCHOUICHA, R., J. WU, AND F. M. VIEIDER (2023): “Choice Lists and “standard Patterns” of Risk-taking,” *Working Paper, Ghent University, Faculty of Economics and Business Administration*.
- BOWERS, J. S. AND C. J. DAVIS (2012): “Bayesian just-so stories in psychology and neuroscience.” *Psychological bulletin*, 138, 389.
- BRADLEY, S. AND N. E. FELDMAN (2020): “Hidden baggage: Behavioral responses to changes in airline ticket tax disclosure,” *American Economic Journal: Economic Policy*, 12, 58–87.
- BRONCHETTI, E. T., J. B. KESSLER, E. B. MAGENHEIM, D. TAUBINSKY, AND E. ZWICK (2023): “Is Attention Produced Optimally? Theory and Evidence From Experiments With Bandwidth Enhancements,” *Econometrica*, 91, 669–707.
- BROWN, J., T. HOSSAIN, AND J. MORGAN (2010): “Shrouded attributes and information suppression: Evidence from the field,” *The Quarterly Journal of Economics*, 125, 859–876.
- BRUHIN, A., M. MANAI, AND L. SANTOS-PINTO (2022): “Risk and rationality: The relative importance of probability weighting and choice set dependence,” *Journal of Risk and Uncertainty*, 65, 139–184.
- BURSZTYN, L., A. W. CAPPELEN, B. TUNGODDEN, A. VOENA, AND D. H. YANAGIZAWA-DROTT (2023): “How Are Gender Norms Perceived?” Tech. rep., National Bureau of Economic Research.
- BUSHONG, B., M. RABIN, AND J. SCHWARTZSTEIN (2021): “A model of relative thinking,” *The Review of Economic Studies*, 88, 162–191.

- CALFORD, E. M. AND T. N. CASON (2022): “Contingent Reasoning and Dynamic Public Goods Provision,” .
- CAMARA, M. K. (2022): “Computationally Tractable Choice.” in *EC*, 28.
- CAMERER, C. (2008): “The case for mindful economics,” *Foundations of positive and normative economics*, 43–69.
- CAMERER, C., L. BABCOCK, G. LOEWENSTEIN, AND R. THALER (1997): “Labor supply of New York City cabdrivers: One day at a time,” *The Quarterly Journal of Economics*, 112, 407–441.
- CAMERER, C. AND T. HUA HO (1999): “Experience-weighted attraction learning in normal form games,” *Econometrica*, 67, 827–874.
- CAMERER, C., G. LOEWENSTEIN, AND D. PRELEC (2005): “Neuroeconomics: How neuroscience can inform economics,” *Journal of economic Literature*, 43, 9–64.
- CAMERER, C. F. (1987): “Do biases in probability judgment matter in markets? Experimental evidence,” *The American Economic Review*, 77, 981–997.
- CAPLIN, A., D. CSABA, J. LEAHY, AND O. NOV (2020): “Rational inattention, competitive supply, and psychometrics,” *The Quarterly Journal of Economics*, 135, 1681–1724.
- CAPLIN, A. AND M. DEAN (2015): “Revealed Preference, Rational Inattention, and Costly Information Acquisition,” *American Economic Review*, 105, 2183–2203.
- CARD, D., S. DELLAVIGNA, C. JIANG, AND D. TAUBINSKY (2024): “Understanding Expert Choices Using Decision Time,” Tech. rep., National Bureau of Economic Research.
- CARRERA, M., H. ROYER, M. STEHR, J. SYDNOR, AND D. TAUBINSKY (2022): “Who chooses commitment? Evidence and welfare implications,” *The Review of Economic Studies*, 89, 1205–1244.
- CARVALHO, L. AND D. SILVERMAN (2019): “Complexity and sophistication,” Tech. rep., National Bureau of Economic Research.
- (2024): “Complexity and sophistication,” *Journal of Political Economy Microeconomics*, 2, 43–76.
- CERIGIONI, F. (2021): “Dual decision processes: Retrieving preferences when some choices are automatic,” *Journal of Political Economy*, 129, 1667–1704.
- CERREIA-VIOGLIO, S., D. DILLENBERGER, AND P. ORTOLEVA (2015): “Cautious expected utility and the certainty effect,” *Econometrica*, 83, 693–728.

- (2022): “Caution and Reference Effects,” *Working Paper*.
- CHAKRABORTY, A. (2021): “Present bias,” *Econometrica*, 89, 1921–1961.
- CHAKRABORTY, A., E. M. CALFORD, G. FENIG, AND Y. HALEVY (2017): “External and internal consistency of choices made in convex time budgets,” *Experimental economics*, 20, 687–706.
- CHAKRABORTY, A. AND C. W. KENDALL (2022): “Noisy foresight,” Tech. rep., National Bureau of Economic Research.
- CHANDRASEKHAR, A. G., H. LARREGUY, AND J. P. XANDRI (2020): “Testing models of social learning on networks: Evidence from two experiments,” *Econometrica*, 88, 1–32.
- CHAPMAN, J., M. DEAN, P. ORTOLEVA, E. SNOWBERG, AND C. CAMERER (2023a): “Econographics,” *Journal of Political Economy Microeconomics*, 1, 115–161.
- (2023b): “Willingness to accept, willingness to pay, and loss aversion,” Tech. rep., National Bureau of Economic Research.
- CHARLES, C. (2022a): “Memory and trading,” *Available at SSRN 3759444*.
- (2022b): “Memory Moves Markets,” *Available at SSRN 4019728*.
- CHARLES, C., C. FRYDMAN, AND M. KILIC (2022): “Insensitive Investors,” .
- CHARNESS, G. AND D. LEVIN (2009): “The Origin of the Winner’s Curse: A Laboratory Study,” *American Economic Journal: Microeconomics*, 207–236.
- CHATER, N., M. OAKSFORD, ET AL. (2008): *The probabilistic mind: Prospects for Bayesian cognitive science*, OUP Oxford.
- CHETTY, R., A. LOONEY, AND K. KROFT (2009): “Salience and taxation: Theory and evidence,” *American economic review*, 99, 1145–77.
- CHOI, S., J. KIM, E. LEE, AND J. LEE (2022): “Probability weighting and cognitive ability,” *Management Science*, 68, 5201–5215.
- COFFMAN, K. B. (2014): “Evidence on self-stereotyping and the contribution of ideas,” *The Quarterly Journal of Economics*, 129, 1625–1660.
- COHEN, M., J.-Y. JAFFRAY, AND T. SAID (1987): “Experimental comparison of individual behavior under risk and under uncertainty for gains and for losses,” *Organizational behavior and human decision processes*, 39, 1–22.
- COHEN, S. AND S. LI (2022): “Sequential Cursed Equilibrium,” *arXiv preprint arXiv:2212.06025*.

- CONLISK, J. (1996): “Why bounded rationality?” *Journal of economic literature*, 34, 669–700.
- CONLON, J. J. (2023): “Attention, Neglect, and Persuasion,” .
- CONLON, J. J. AND D. PATEL (2022): “What jobs come to mind? stereotypes about fields of study,” Tech. rep., Working Paper.
- CUBITT, R. P., D. NAVARRO-MARTINEZ, AND C. STARMER (2015): “On preference imprecision,” *Journal of Risk and Uncertainty*, 50, 1–34.
- DA SILVEIRA, R. A., Y. SUNG, AND M. WOODFORD (2020): “Optimally imprecise memory and biased forecasts,” Tech. rep., National Bureau of Economic Research.
- D’ACUNTO, F., D. HOANG, M. PALOVIITA, AND M. WEBER (2019): “IQ, expectations, and choice,” Tech. rep., National Bureau of Economic Research.
- (2023): “Cognitive constraints and economic incentives,” *Bank of Finland Research Discussion Paper*.
- DAL BÓ, E., P. DAL BÓ, AND E. EYSTER (2018): “The demand for bad policy when voters underappreciate equilibrium effects,” *The Review of Economic Studies*, 85, 964–998.
- DANZ, D., L. VESTERLUND, AND A. J. WILSON (2022): “Belief elicitation and behavioral incentive compatibility,” *American Economic Review*, 112, 2851–2883.
- DE CLIPPEL, G. AND K. ROZEN (2023): “Bounded Rationality in Choice Theory: A Survey,” *Working Paper*.
- DEAN, M. AND N. N. NELIGH (2023): “Experimental tests of rational inattention,” *Journal of Political Economy*.
- DEAN, M. AND P. ORTOLEVA (2019): “The empirical relationship between nonstandard economic behaviors,” *Proceedings of the National Academy of Sciences*, 116, 16262–16267.
- DECLIPPEL, G., S. KARIV, R. OPREA, AND K. ROZEN (2024): “As If,” *Working Paper*.
- DELLAVIGNA, S. AND M. GENTZKOW (2019): “Uniform pricing in us retail chains,” *The Quarterly Journal of Economics*, 134, 2011–2084.
- DERTWINKEL-KALT, M., H. GERHARDT, G. RIENER, F. SCHWERTER, AND L. STRANG (2022): “Concentration bias in intertemporal choice,” *The Review of Economic Studies*, 89, 1314–1334.
- DERTWINKEL-KALT, M., K. KÖHLER, M. R. LANGE, AND T. WENZEL (2017): “Demand shifts due to salience effects: Experimental evidence,” *Journal of the European Economic Association*, 15, 626–653.

- DETRWINKEL-KALT, M. AND M. KÖSTER (2020a): “Salience and skewness preferences,” *Journal of the European Economic Association*, 18, 2057–2107.
- (2020b): “Salient Cues and Complexity,” *Available at SSRN 3697313*.
- DIAMOND, P. A. AND J. A. HAUSMAN (1994): “Contingent valuation: is some number better than no number?” *Journal of economic perspectives*, 8, 45–64.
- DIMMOCK, S. G., R. KOUWENBERG, AND P. P. WAKKER (2015): “Ambiguity attitudes in a large representative sample,” *Management Science*, 62, 1363–1380.
- DOHMEN, T., A. FALK, D. HUFFMAN, AND U. SUNDE (2010): “Are risk aversion and impatience related to cognitive ability?” *American Economic Review*, 100, 1238–60.
- DRERUP, T., B. ENKE, AND H.-M. VON GAUDECKER (2017): “The precision of subjective data and the explanatory power of economic models,” *Journal of Econometrics*, 200, 378–389.
- DRUGOWITSCH, J., V. WYART, A.-D. DEVAUCHELLE, AND E. KOECHLIN (2016): “Computational precision of mental inference as critical source of human choice suboptimality,” *Neuron*, 92, 1398–1411.
- EBERT, J. E. (2001): “The role of cognitive resources in the valuation of near and far future events,” *Acta psychologica*, 108, 155–171.
- EBERT, J. E. AND D. PRELEC (2007): “The fragility of time: Time-insensitivity and valuation of the near and far future,” *Management science*, 53, 1423–1438.
- ELLIS, A. AND D. J. FREEMAN (2020): “Revealing choice bracketing,” *arXiv preprint arXiv:2006.14869*.
- ELLIS, A. AND Y. MASATLIOGLU (2019): “A behavioral foundation for endogenous salience,” .
- ENGELMANN, J., A. HIRMAS, AND J. J. VAN DER WEELE (2021): “Top down or bottom up? disentangling the channels of attention in risky choice,” .
- ENKE, B. (2020): “What you see is all there is,” *The Quarterly Journal of Economics*, 135, 1363–1398.
- ENKE, B. AND T. GRAEBER (2023): “Cognitive uncertainty,” *Quarterly Journal of Economics*.
- (2024): “Tradeoff Uncertainty and Comparative Thinking,” *Working Paper*.
- ENKE, B., T. GRAEBER, AND R. OPREA (2023a): “Complexity and Hyperbolic Discounting,” Tech. rep., National Bureau of Economic Research.

- (2023b): “Confidence, Self-Selection and Bias in the Aggregate,” *American Economic Review*.
- ENKE, B., T. GRAEBER, R. OPREA, AND J. YANG (2024a): “Behavioral Attenuation,” *Working Paper*.
- ENKE, B., F. SCHWERTER, AND F. ZIMMERMANN (2024b): “Associative Memory, Beliefs and Market Interactions,” *Journal of Financial Economics*.
- ENKE, B. AND C. SHUBATT (2023): “Quantifying Lottery Choice Complexity,” *Working Paper*.
- ENKE, B. AND F. ZIMMERMANN (2019): “Correlation Neglect in Belief Formation,” *Review of Economic Studies*, 86, 313–332.
- EREV, I., E. ERT, A. E. ROTH, E. HARUVY, S. M. HERZOG, R. HAU, R. HERTWIG, T. STEWART, R. WEST, AND C. LEBIERE (2010): “A choice prediction competition: Choices from experience and from description,” *Journal of Behavioral Decision Making*, 23, 15–47.
- ESPONDA, I., R. OPREA, AND S. YUKSEL (2023): “Seeing what is representative,” *The Quarterly Journal of Economics*, qjad020.
- ESPONDA, I. AND D. POUZO (2016): “Berk–Nash equilibrium: A framework for modeling agents with misspecified models,” *Econometrica*, 84, 1093–1130.
- ESPONDA, I. AND E. VESPA (2014): “Hypothetical Thinking and Information Extraction in the Laboratory,” *American Economic Journal: Microeconomics*, 6, 180–202.
- (2023): “Contingent thinking and the sure-thing principle: Revisiting classic anomalies in the laboratory,” *Review of Economic Studies*.
- ESPONDA, I., E. VESPA, AND S. YUKSEL (2020): “Mental models and learning: The case of base-rate neglect,” Tech. rep., Tech. rep.
- EVERS, E. R., A. IMAS, AND C. KANG (2022): “On the role of similarity in mental accounting and hedonic editing,” *Psychological review*, 129, 777.
- EXLEY, C. L. AND J. B. KESSLER (2019): “Motivated errors,” Tech. rep., National Bureau of Economic Research.
- (2021): “Information avoidance and image concerns,” Tech. rep., National Bureau of Economic Research.
- EYSTER, E. AND M. RABIN (2005): “Cursed equilibrium,” *Econometrica*, 73, 1623–1672.
- FALK, A., A. BECKER, T. DOHMEN, B. ENKE, D. HUFFMAN, AND U. SUNDE (2018): “Global evidence on economic preferences,” *The Quarterly Journal of Economics*, 133, 1645–1692.

- FAN, T., Y. LIANG, AND C. PENG (2023): “The inference-forecast gap in belief updating,” *Available at SSRN 3889069*.
- FEHR, E. AND K. WU (2023): “Obfuscation in competitive markets,” Tech. rep., Working Paper.
- FIEDLER, K. AND P. JUSLIN (2006): *Information sampling and adaptive cognition*, Cambridge University Press.
- FISCHHOFF, B. AND W. BRUINE DE BRUIN (1999): “Fifty–fifty= 50%?” *Journal of Behavioral Decision Making*, 12, 149–163.
- FRÉCHETTE, G. R., A. LIZZERI, AND J. PEREGO (2022): “Rules and commitment in communication: An experimental analysis,” *Econometrica*, 90, 2283–2318.
- FREDERICK, S. (2005): “Cognitive Reflection and Decision Making,” *Journal of Economic Perspectives*, 19, 25–42.
- FRIEDMAN, D. (1989): “The S-shaped value function as a constrained optimum,” *The American Economic Review*, 79, 1243–1248.
- FRYDMAN, C. AND L. J. JIN (2022): “Efficient coding and risky choice,” *The Quarterly Journal of Economics*, 137, 161–213.
- (2023): “On the Source and Instability of Probability Weighting,” .
- FRYDMAN, C. AND M. M. MORMANN (2018): “The Role of Saliency in Choice Under Risk: An Experimental Investigation,” *Working Paper*.
- FRYDMAN, C. AND S. NUNNARI (2023): “Coordination with Cognitive Noise,” *Working Paper*.
- FRYER, R. AND M. O. JACKSON (2008): “A categorical model of cognition and biased decision making,” *The BE Journal of Theoretical Economics*, 8, 0000102202193517041357.
- FUDENBERG, D. (2006): “Advancing Beyond “Advances in Behavioral Economics”,” *Journal of Economic Literature*, 44, 694–711.
- FUDENBERG, D. AND G. LANZANI (2023): “Which misperceptions persist?” Tech. rep.
- FUDENBERG, D., G. LANZANI, AND P. STRACK (2022): “Selective memory equilibrium,” *Available at SSRN 4015313*.
- FUDENBERG, D. AND D. K. LEVINE (2006): “A dual-self model of impulse control,” *American economic review*, 96, 1449–1476.
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2018): “Speed, accuracy, and the optimal timing of choices,” *American Economic Review*, 108, 3651–3684.

- GABAIX, X. (2014): “A sparsity-based model of bounded rationality,” *The Quarterly Journal of Economics*, 129, 1661–1710.
- (2019): “Behavioral inattention,” in *Handbook of behavioral economics: Applications and foundations 1*, Elsevier, vol. 2, 261–343.
- (2023): “Behavioral macroeconomics via sparse dynamic programming,” *Journal of the European Economic Association*, jvad057.
- GABAIX, X. AND T. GRAEBER (2023): “The Complexity of Economic Decisions: Theory and Evidence,” *Working Paper*.
- GABAIX, X. AND D. LAIBSON (2000): “A boundedly rational decision algorithm,” *American Economic Review*, 90, 433–438.
- (2005): “Bounded rationality and directed cognition,” *Harvard University*, 20.
- (2022): “Myopia and discounting,” Tech. rep., National bureau of economic research.
- GABAIX, X., D. LAIBSON, G. MOLOCHE, AND S. WEINBERG (2011): “Costly information acquisition: Experimental analysis of a boundedly rational model,” *American Economic Review*, 96, 2899–2922.
- GAGNON-BARTSCH, T., M. RABIN, AND J. SCHWARTZSTEIN (2021): “Channeled attention and stable errors,” *Working Paper*.
- GENNAIOLI, N. AND A. SHLEIFER (2010): “What Comes to Mind,” *Quarterly Journal of Economics*, 125, 1399–1433.
- GERHARDT, H., G. P. BIELE, H. R. HEEKEREN, AND H. UHLIG (2016): “Cognitive load increases risk aversion,” *Working Paper*.
- GERSHMAN, S. (2021): *What makes us smart: The computational logic of human cognition*, Princeton University Press.
- GERSHMAN, S. J. AND R. BHUI (2019): “Rationally inattentive intertemporal choice,” *bioRxiv*, 680652.
- GIGERENZER, G. AND W. GAISSMAIER (2011): “Heuristic decision making,” *Annual review of psychology*, 62, 451–482.
- GIGERENZER, G. AND P. M. TODD (1999): *Simple heuristics that make us smart*, Oxford University Press, USA.
- GIGLIO, S., M. MAGGIORI, J. STROEBEL, AND S. UTKUS (2021): “Five facts about beliefs and portfolios,” *American Economic Review*, 111, 1481–1522.

- GILBOA, I. AND D. SCHMEIDLER (1995): “Case-based decision theory,” *The quarterly Journal of economics*, 110, 605–639.
- GILLEN, B., E. SNOWBERG, AND L. YARIV (2019): “Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study,” *Journal of Political Economy*.
- GIUSTINELLI, P., C. F. MANSKI, AND F. MOLINARI (2019): “Precise or Imprecise Probabilities? Evidence from Survey Response on Late-onset Dementia,” Tech. rep., National Bureau of Economic Research.
- GRAEBER, T. (2022): “Inattentive Inference,” *Journal of the European Economic Association*.
- GRAEBER, T., S. NOY, AND C. ROTH (2023a): “Lost in transmission,” *Working Paper*.
- GRAEBER, T., C. ROTH, AND C. SCHESCH (2023b): “Contagious Beliefs,” *Working Paper*.
- GRAEBER, T., C. ROTH, AND F. ZIMMERMANN (2022): “Stories, Statistics and Memory,” *Working Paper*.
- GRIFFITHS, T., C. KEMP, AND J. TENENBAUM (2008): “Bayesian models of cognition,” *Cambridge Handbook of Computational Cognitive Modeling*.
- GRIMM, V. AND F. MENGEL (2014): “An Experiment on Belief Formation in Networks,” *Working Paper*.
- GUL, F. (1991): “A theory of disappointment aversion,” *Econometrica: Journal of the Econometric Society*, 667–686.
- HAAN, P., C. SUN, U. SUNDE, AND G. WEIZSACKER (2022): “Non-additivity of subjective expectations over different time intervals,” .
- HAISLEY, E. C. AND R. A. WEBER (2010): “Self-serving interpretations of ambiguity in other-regarding behavior,” *Games and economic behavior*, 68, 614–625.
- HALEVY, Y., D. WALKER-JONES, AND L. ZRILL (2023): “Difficult Decisions,” Tech. rep.
- HANDEL, B. AND J. SCHWARTZSTEIN (2018): “Frictions or mental gaps: what’s behind the information we (don’t) use and when do we care?” *Journal of Economic Perspectives*, 32, 155–178.
- HANNA, R., S. MULLAINATHAN, AND J. SCHWARTZSTEIN (2014): “Learning Through Noticing: Theory and Experimental Evidence in Farming,” *Quarterly Journal of Economics*, 129, 1311–1353.
- HARBAUGH, W. T., K. KRAUSE, AND L. VESTERLUND (2010): “The fourfold pattern of risk attitudes in choice and pricing tasks,” *The Economic Journal*, 120, 595–611.

- HARSTAD, R. M. AND R. SELTEN (2013): “Bounded-rationality models: tasks to become intellectually competitive,” *Journal of Economic Literature*, 51, 496–511.
- HARTZMARK, S. M., S. D. HIRSHMAN, AND A. IMAS (2021): “Ownership, learning, and beliefs,” *The Quarterly journal of economics*, 136, 1665–1717.
- HARTZMARK, S. M. AND K. SHUE (2018): “A tough act to follow: Contrast effects in financial markets,” *The Journal of Finance*, 73, 1567–1613.
- HASTINGS, J. S. AND J. M. SHAPIRO (2013): “Fungibility and consumer choice: Evidence from commodity price shocks,” *The quarterly journal of economics*, 128, 1449–1498.
- HAZELL, J., C. PATTERSON, H. SARSONS, AND B. TASKA (2022): “National wage setting,” Tech. rep., National Bureau of Economic Research.
- HE, J. AND P. NATENZON (2022): “Moderate utility,” *Available at SSRN 3243657*.
- (2023): “Random Choice and Differentiation,” *Working Paper*.
- HE, L., R. GOLMAN, AND S. BHATIA (2019): “Variable time preference,” *Cognitive psychology*, 111, 53–79.
- HEIDHUES, P., B. KŐSZEGI, AND P. STRACK (2018): “Unrealistic expectations and misguided learning,” *Econometrica*, 86, 1159–1214.
- HENKEL, L. (2022): “Experimental evidence on the relationship between perceived ambiguity and likelihood insensitivity,” Tech. rep., ECONtribute Discussion Paper.
- HEROLD, F. AND N. NETZER (2023): “Second-best probability weighting,” *Games and Economic Behavior*, 138, 112–125.
- HERWEG, F. AND D. MÜLLER (2021): “A comparison of regret theory and salience theory for decisions under risk,” *Journal of Economic Theory*, 193, 105226.
- HSEE, C. K. (1996): “The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives,” *Organizational behavior and human decision processes*, 67, 247–257.
- HSEE, C. K., G. F. LOEWENSTEIN, S. BLOUNT, AND M. H. BAZERMAN (1999): “Preference reversals between joint and separate evaluations of options: A review and theoretical analysis.” *Psychological bulletin*, 125, 576.
- HU, E. H. (2023): “A Procedural Model of Complexity Under Risk,” .

- HUCK, S. AND G. WEIZSÄCKER (1999): “Risk, complexity, and deviations from expected-value maximization: Results of a lottery choice experiment,” *Journal of Economic Psychology*, 20, 699–715.
- HVIDBERG, K. B., C. T. KREINER, AND S. STANTCHEVA (2023): “Social Positions and Fairness Views on Inequality,” *Review of Economic Studies*.
- ILUT, C. AND R. VALCHEV (2023): “Economic agents as imperfect problem solvers,” *The Quarterly Journal of Economics*, 138, 313–362.
- IMAS, A., M. H. JUNG, S. SACCARDO, AND J. VOSGERAU (2022): “The Impact of Joint versus Separate Prediction Mode on Forecasting Accuracy,” Tech. rep., National Bureau of Economic Research.
- IMAS, A., M. KUHN, AND V. MIRONOVA (2021): “Waiting to choose: The Role of Deliberation in Intertemporal Choice,” *American Economic Journal: Microeconomics*.
- IYENGAR, S. S. AND E. KAMENICA (2010): “Choice proliferation, simplicity seeking, and asset allocation,” *Journal of Public Economics*, 94, 530–539.
- JEHIEL, P. (2005): “Analogy-based expectation equilibrium,” *Journal of Economic theory*, 123, 81–104.
- JIANG, Z., H. LIU, C. PENG, AND H. YAN (2022): “Investor Memory and Biased Beliefs: Evidence from the Field,” *Available at SSRN*.
- JIN, G. Z., M. LUCA, AND D. MARTIN (2021): “Is no news (perceived as) bad news? An experimental investigation of information disclosure,” *American Economic Journal: Microeconomics*, 13, 141–173.
- (2022): “Complex disclosure,” *Management Science*, 68, 3236–3261.
- JOHNSON, E. J., C. CAMERER, S. SEN, AND T. RYMON (2002): “Detecting failures of backward induction: Monitoring information search in sequential bargaining,” *Journal of economic theory*, 104, 16–47.
- KAATS, M. (2023): “Valuing Health,” *Working Paper*.
- KAGEL, J. H. (1990): *Auctions: a survey of experimental research*, University of Pittsburgh.
- KAHANA, M. J. (2012): *Foundations of human memory*, OUP USA.
- KAHNEMAN, D. (2011): *Thinking, Fast and Slow*, Macmillan.

- KAHNEMAN, D., S. FREDERICK, ET AL. (2002): “Representativeness revisited: Attribute substitution in intuitive judgment,” *Heuristics and biases: The psychology of intuitive judgment*, 49, 74.
- KAHNEMAN, D. AND A. TVERSKY (1972): “Subjective probability: A judgment of representativeness,” *Cognitive psychology*, 3, 430–454.
- (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–292.
- KAMENICA, E. (2008): “Contextual inference in markets: On the informational content of product lines,” *American Economic Review*, 98, 2127–49.
- KENDALL, C. AND A. CHAKRABORTY (2022): “Future self-proof elicitation mechanisms,” *Available at SSRN 4032946*.
- KENDALL, C. AND R. OPREA (2022): “On the Complexity of Forming Mental Models,” *Working Paper*.
- KENDALL, C. W. AND C. CHARLES (2022): “Causal narratives,” Tech. rep., National Bureau of Economic Research.
- KHAW, M. W., Z. LI, AND M. WOODFORD (2021): “Cognitive imprecision and small-stakes risk aversion,” *The review of economic studies*, 88, 1979–2013.
- (2022): “Cognitive Imprecision and Stake-Dependent Risk Attitudes,” Tech. rep., National Bureau of Economic Research.
- KŐSZEGI, B., G. LOEWENSTEIN, AND T. MUROOKA (2022): “Fragile Self-Esteem,” *The Review of Economic Studies*, 89, 2026–2060.
- KŐSZEGI, B. AND F. MATĚJKA (2020): “Choice simplification: A theory of mental budgeting and naive diversification,” *The Quarterly Journal of Economics*, 135, 1153–1207.
- KŐSZEGI, B. AND M. RABIN (2006): “A model of reference-dependent preferences,” *The Quarterly Journal of Economics*, 121, 1133–1165.
- KŐSZEGI, B. AND A. SZEIDL (2013): “A model of focusing in economic choice,” *Quarterly Journal of Economics*, 128, 53–104.
- KRAJBICH, I., D. LU, C. CAMERER, AND A. RANGEL (2012): “The attentional drift-diffusion model extends to simple purchasing decisions,” *Frontiers in psychology*, 3, 193.
- LAIBSON, D. (2001): “A cue-theory of consumption,” *The Quarterly Journal of Economics*, 116, 81–119.

- LANDRY, P. AND R. WEBB (2021): “Pairwise normalization: A neuroeconomic theory of multi-attribute choice,” *Journal of Economic Theory*, 193, 105221.
- LANZANI, G. (2022a): “Correlation made simple: Applications to salience and regret theory,” *The Quarterly Journal of Economics*, 137, 959–987.
- (2022b): “Dynamic Concern for Misspecification,” *Available at SSRN 4454504*.
- LI, S. (2017): “Obviously strategy-proof mechanisms,” *American Economic Review*, 107, 3257–3287.
- (2024): “Designing Simple Mechanisms,” *arXiv preprint arXiv:2403.18694*.
- LI, X. AND C. F. CAMERER (2022): “Predictable Effects of Visual Salience in Experimental Decisions and Games,” *The Quarterly Journal of Economics*, 137, 1849–1900.
- LIAN, C. (2021): “A theory of narrow thinking,” *The Review of Economic Studies*, 88, 2344–2374.
- LIANG, Y. (2023a): “Boundedly Rational Information Demand,” .
- (2023b): “Learning from Unknown Information Sources,” *Management Science*.
- LIEDER, F. AND T. L. GRIFFITHS (2020): “Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources,” *Behavioral and brain sciences*, 43.
- LIU, S. AND N. NETZER (2023): “Happy Times: Measuring Happiness Using Response Times,” .
- LOEWENSTEIN, G. AND D. PRELEC (1992): “Anomalies in intertemporal choice: Evidence and an interpretation,” *The Quarterly Journal of Economics*, 107, 573–597.
- LOEWENSTEIN, G. AND Z. WOJTOWICZ (2023): “The Economics of Attention,” *Available at SSRN 4368304*.
- LOOMES, G. AND R. SUGDEN (1982): “Regret theory: An alternative theory of rational choice under uncertainty,” *The economic journal*, 92, 805–824.
- (1995): “Incorporating a stochastic element into decision theories,” *European Economic Review*, 39, 641–648.
- LUTTMER, E. F., P. DE OLIVEIRA, AND T. DMITRY (2023): “Failures of Contingent Reasoning in Annuity Decisions,” *Working Paper*.
- MAĆKOWIAK, B., F. MATĚJKA, AND M. WIEDERHOLT (2023): “Rational inattention: A review,” *Journal of Economic Literature*, 61, 226–273.

- MALMENDIER, U. (2021): “Experience effects in finance: Foundations, applications, and future directions,” *Review of Finance*, 25, 1339–1363.
- MALMENDIER, U. AND S. NAGEL (2011): “Depression babies: do macroeconomic experiences affect risk taking?” *The quarterly journal of economics*, 126, 373–416.
- (2016): “Learning from inflation experiences,” *The Quarterly Journal of Economics*, 131, 53–87.
- MALMENDIER, U., S. NAGEL, AND Z. YAN (2021): “The making of hawks and doves,” *Journal of Monetary Economics*, 117, 19–42.
- MALMENDIER, U. AND L. VELDKAMP (2022): “Information resonance,” Tech. rep., Working Paper.
- MALMENDIER, U. AND J. A. WACHTER (2021): “Memory of past experiences and economic decisions,” *Available at SSRN 4013583*.
- MARTÍNEZ-MARQUINA, A., M. NIEDERLE, AND E. VESPA (2019): “Failures in contingent reasoning: The role of uncertainty,” *American Economic Review*, 109, 3437–74.
- MASATLIOGLU, Y. AND E. A. OK (2005): “Rational choice with status quo bias,” *Journal of economic theory*, 121, 1–29.
- MASSEY, C. AND G. WU (2005): “Detecting regime shifts: The causes of under-and overreaction,” *Management Science*, 51, 932–947.
- MCGRANAGHAN, C., K. NIELSEN, T. O’DONOGHUE, J. SOMERVILLE, AND C. SPRENGER (2022): “Distinguishing common ratio preferences from common ratio effects using paired valuation tasks,” *Unpublished Manuscript*.
- MEDIN, D. L., R. L. GOLDSTONE, AND D. GENTNER (1993): “Respects for similarity,” *Psychological review*, 100, 254.
- MOORE, D. A. AND P. J. HEALY (2008): “The trouble with overconfidence,” *Psychological review*, 115, 502.
- MORRISON, W. AND D. TAUBINSKY (2019): “Rules of thumb and attention elasticities: Evidence from under-and overreaction to taxes,” *The Review of Economics and Statistics*, 1–45.
- MULLAINATHAN, S. (2002): “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 735–774.
- MULLAINATHAN, S. AND Z. OBERMEYER (2022): “Diagnosing physician error: A machine learning approach to low-value health care,” *The Quarterly Journal of Economics*, 137, 679–727.

- MULLAINATHAN, S., J. SCHWARTZSTEIN, AND A. SHLEIFER (2008): “Coarse thinking and persuasion,” *The Quarterly journal of economics*, 123, 577–619.
- NAGEL, L., M. NIEDERLE, AND E. VESPA (2023): “Decomposing the winner’s curse,” *Working Paper*.
- NATENZON, P. (2019): “Random choice and learning,” *Journal of Political Economy*, 127, 419–457.
- NETZER, N. (2009): “Evolution of time preferences and attitudes toward risk,” *American Economic Review*, 99, 937–955.
- NETZER, N., A. J. ROBSON, J. STEINER, AND P. KOCOUREK (2022): “Endogenous risk attitudes,” .
- NGANGOUÉ, M. K. AND G. WEIZSÄCKER (2021): “Learning from unrealized versus realized prices,” *American Economic Journal: Microeconomics*, 13, 174–201.
- NIEDERLE, M. AND E. VESPA (2023): “Cognitive limitations: Failures of contingent thinking,” *Annual Review of Economics*.
- NIELSEN, K. AND J. REHBECK (2022): “When Choices are mistakes,” *American Economic Review*, 112, 2237–68.
- NIELSEN, K. AND L. RIGOTTI (2022): “Revealed Incomplete Preferences,” *arXiv preprint arXiv:2205.08584*.
- OAKSFORD, M. AND N. CHATER (2007): *Bayesian rationality: The probabilistic approach to human reasoning*, Oxford University Press.
- O’DONOGHUE, T. AND C. SPRENGER (2018): “Reference-dependent preferences,” in *Handbook of behavioral economics: Applications and foundations 1*, Elsevier, vol. 1, 1–77.
- OK, E. A., P. ORTOLEVA, AND G. RIELLA (2012): “Incomplete preferences under uncertainty: Indecisiveness in beliefs versus tastes,” *Econometrica*, 80, 1791–1808.
- (2015): “Revealed (p) reference theory,” *American Economic Review*, 105, 299–321.
- OPREA, R. (2020): “What makes a rule complex?” *American economic review*, 110, 3913–51.
- (2022): “Simplicity Equivalents,” *Working Paper*.
- PETZSCHNER, F. H., S. GLASAUER, AND K. E. STEPHAN (2015): “A Bayesian perspective on magnitude estimation,” *Trends in cognitive sciences*, 19, 285–293.
- PRAT-CARRABIN, A. AND M. WOODFORD (2022): “Imprecise Probabilistic Inference from Sequential Data,” .

- PRELEC, D. AND G. LOEWENSTEIN (1991): “Decision making over time and under uncertainty: A common approach,” *Management science*, 37, 770–786.
- PURI, I. (2022): “Simplicity and Risk,” *Available at SSRN 3253494*.
- PYCIA, M. AND P. TROYAN (2023): “A theory of simplicity in games and mechanism design,” *Econometrica*, 91, 1495–1526.
- RABIN, M. (2000): “Risk Aversion and Expected-Utility Theory: A Calibration Theorem,” *Econometrica*, 68, 1281–1292.
- (2013): “Incorporating limited rationality into economics,” *Journal of Economic Literature*, 51, 528–43.
- RABIN, M. AND G. WEIZSÄCKER (2009): “Narrow bracketing and Dominated Choices,” *American Economic Review*, 99, 1508–1543.
- RADBRUCH, J. AND A. SCHIPROWSKI (2022): “Interview sequences and the formation of subjective assessments,” .
- RATCLIFF, R. (1978): “A theory of memory retrieval.” *Psychological review*, 85, 59.
- RATCLIFF, R. AND G. MCKOON (2008): “The diffusion decision model: theory and data for two-choice decision tasks,” *Neural computation*, 20, 873–922.
- READ, D., G. LOEWENSTEIN, AND M. RABIN (1999): “Choice Bracketing,” *Journal of Risk and Uncertainty*, 171–197.
- REES-JONES, A. AND D. TAUBINSKY (2020): “Measuring “schmeduling,”” *The Review of Economic Studies*, 87, 2399–2438.
- REUTSKAJA, E., R. NAGEL, C. F. CAMERER, AND A. RANGEL (2011): “Search dynamics in consumer choice under time pressure: An eye-tracking study,” *American Economic Review*, 101, 900–926.
- ROBSON, A. J. (2001): “The biological basis of economic behavior,” *Journal of economic literature*, 39, 11–33.
- ROZEN, K. (2010): “Foundations of intrinsic habit formation,” *Econometrica*, 78, 1341–1373.
- RUBINSTEIN, A. (1988): “Similarity and decision-making under risk (Is there a utility theory resolution to the Allais paradox?),” *Journal of economic theory*, 46, 145–153.
- (1998): *Modeling bounded rationality*, MIT press.
- (2003): ““Economics and psychology”? The case of hyperbolic discounting,” *International Economic Review*, 44, 1207–1216.

- (2007): “Instinctive and Cognitive Reasoning: A Study of Response Times,” *Economic Journal*, 117, 1243–1259.
- (2016): “A Typology of Players: Between Instinctive and Contemplative,” *Quarterly Journal of Economics*, 131, 859–890.
- SALANT, Y. AND J. L. SPENKUCH (2022): “Complexity and Choice,” Tech. rep., National Bureau of Economic Research.
- SANJURJO, A. (2023): “Complexity in Multiattribute Choice,” *Available at SSRN 4580158*.
- SCHUMACHER, H., I. KESTERNICH, M. KOSFELD, AND J. WINTER (2017): “One, two, many-insensitivity to group size in games with concentrated benefits and dispersed costs,” *The Review of Economic Studies*, 84, 1346–1377.
- SCHWARTZSTEIN, J. (2014): “Selective Attention and Learning,” *Journal of the European Economic Association*, 12, 1423–1452.
- SCHWARTZSTEIN, J. AND A. SUNDERAM (2021): “Using models to persuade,” *American Economic Review*, 111, 276–323.
- (2022): “Shared models in networks, organizations, and groups,” Tech. rep., National Bureau of Economic Research.
- SCHWEITZER, M. E. AND G. P. CACHON (2000): “Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence,” *Management science*, 46, 404–420.
- SHUBATT, C. AND J. YANG (2024): “Similarity and Comparison Complexity,” *Working Paper*.
- SIMON, H. A. (1956): “Rational choice and the structure of the environment,” *Psychological Review*, 63, 129–138.
- (1982): *Models of bounded rationality*, MIT Press.
- SIMS, C. A. (1998): “Stickiness,” in *Carnegie-rochester conference series on public policy*, Elsevier, vol. 49, 317–356.
- (2003): “Implications of rational inattention,” *Journal of monetary Economics*, 50, 665–690.
- SOLTANI, A., B. DE MARTINO, AND C. CAMERER (2012): “A range-normalization model of context-dependent choice: a new model and evidence,” *PLoS computational biology*, 8, e1002607.

- SOMERVILLE, J. (2022): “Range-Dependent Attribute Weighting in Consumer Choice: An Experimental Test,” *Econometrica*, 90, 799–830.
- SONSINO, D., U. BENZION, AND G. MADOR (2002): “The complexity effects on choice with uncertainty—Experimental evidence,” *The Economic Journal*, 112, 936–965.
- SPIEGLER, R. (2016): “Bayesian networks and boundedly rational expectations,” *The Quarterly Journal of Economics*, 131, 1243–1290.
- (2019): “Behavioral economics and the atheoretical style,” *American Economic Journal: Microeconomics*, 11, 173–94.
- (2020): “Behavioral implications of causal misperceptions,” *Annual Review of Economics*, 12, 81–106.
- STANGO, V. AND J. ZINMAN (2023): “We Are All Behavioural, More, or Less: A Taxonomy of Consumer Decision-Making,” *The Review of Economic Studies*, 90, 1470–1498.
- STEINER, J. AND C. STEWART (2016): “Perceiving prospects properly,” *American Economic Review*, 106, 1601–31.
- STRACK, P. AND D. TAUBINSKY (2021): “Dynamic preference “reversals” and time inconsistency,” Tech. rep., National Bureau of Economic Research.
- TOMA, M. AND E. BELL (2022): “Understanding and Increasing Policymakers’ Sensitivity to Program Impact,” *Working Paper*.
- TRAUTMANN, S. T. AND G. VAN DE KUILEN (2015): “Ambiguity attitudes,” *The Wiley Blackwell handbook of judgment and decision making*, 1, 89–116.
- TVERSKY, A. (1969): “Intransitivity of preferences.” *Psychological review*, 76, 31.
- (1977): “Features of similarity.” *Psychological review*, 84, 327.
- TVERSKY, A. AND E. SHAFIR (1992): “Choice under conflict: The dynamics of deferred decision,” *Psychological science*, 3, 358–361.
- VIEIDER, F. M. (2018): “Violence and risk preference: experimental evidence from Afghanistan: comment,” *American Economic Review*, 108, 2366–82.
- (2021): “Noisy coding of time and reward discounting,” Tech. rep., Ghent University, Faculty of Economics and Business Administration.
- (2022): “Decisions under Uncertainty as Bayesian Inference on Choice Options,” *Working Paper*.

- WACHTER, J. A. AND M. J. KAHANA (2019): “A retrieved-context theory of financial decisions,” Tech. rep., National Bureau of Economic Research.
- (2023): “Associative learning and representativeness,” *Available at SSRN 3602327*.
- WALASEK, L. AND G. D. BROWN (2021): “Incomparability and incommensurability in choice: No common currency of value,” *Perspectives on Psychological Science*.
- WEBB, R., P. W. GLIMCHER, AND K. LOUIE (2021): “The normalization of consumer valuations: Context-dependent preferences from neurobiological constraints,” *Management Science*, 67, 93–125.
- WEBB, R., C. ZHAO, M. OSBORNE, P. LANDRY, AND C. CAMERER (2022): “A neuro-autopilot theory of habit: Evidence from canned tuna,” *Rotman School of Management Working Paper*.
- WILSON, A. (2014): “Bounded memory and biases in information processing,” *Econometrica*, 82, 2257–2294.
- WOODFORD, M. (2012a): “Inattentive valuation and reference-dependent choice,” .
- (2012b): “Prospect theory as efficient perceptual distortion,” *American Economic Review*, 102, 41–46.
- (2020): “Modeling imprecision in perception, valuation, and choice,” *Annual Review of Economics*, 12, 579–601.
- XIANG, Y., T. GRAEBER, B. ENKE, AND S. GERSHMAN (2021): “Confidence and Central Tendency in Perceptual Judgment,” *Attention, Perception and Psychophysics*.
- YANG, J. (2023): “On the Decision-Relevance of Subjective Beliefs,” *Working Paper*.