Memory, Attention, and Choice

Pedro Bordalo, Nicola Gennaioli, Andrei Shleifer[1]

February 4, 2019

1. Introduction

Memory appears to play a central role even in simplest choices. Consider a consumer deciding whether to buy an orange at a cafeteria for $1. In the neoclassical model, this consumer would be assumed to know exactly whether that orange is worth more than $1 to him, and make the decision accordingly. In reality, the consumer might need to figure that out. He would spontaneously recall other instances of having an orange, and other instances of spending money on oranges, and estimate based on this recollection whether the orange is worth $1. The more he attends to thoughts about the pleasure of eating an orange, the more likely he is to buy. The more he attends to thoughts about the pain of paying, the less likely he is to buy (Hare, Malmaud, Rangel 2011). Even a monkey choosing between an orange and a banana in an experiment presumably relies on its memory of eating fruit.

But how does memory shape this estimation of value? Can we use the available research on memory to improve our understanding of choice, and help come to grips with puzzling evidence on human decisions? And how does memory interact with attention in shaping choice? Can modeling memory and attention help account for and unify various findings and anomalies of choice, as well as make new predictions?

This paper explores the possibilities of explicitly incorporating empirically established features of human memory into models of economic choice. Evidence from behavioral economics points to two pertinent mechanisms. First, memory is an anchor for valuation. When thinking of a possible stock market investment, we draw on memories of past experienced returns (Malmendier and Nagel 2011). When deciding whether to buy a convertible car, we are influenced by the highly available weather on the day of

[1] The authors are from University of Oxford, Universita Bocconi, and Harvard University, respectively. The present paper replaces, and completely reworks, a 2017 manuscript with the same title. Shleifer thanks the Sloan Foundation and the Pershing Square Venture Fund for Research on the Foundations of Human Behavior for financial support.

the purchase (Busse et al 2015). When choosing a product in a supermarket aisle, some price components are not visible, and we draw – or fail to draw -- upon memory to assess them (Chetty et al 2009). In many economic decisions, memory is thus essential to establish valuations and make choices (Gilboa and Schmeidler 1995, Mullainathan 2002). Yet, as the examples in this paragraph illustrate, its selective nature may sometimes cause problems. It may lead us to neglect informative past returns that we did not experience, weather conditions that are not available today, or sales taxes paid at the checkout counter.

Because memory anchors valuation, it also creates a reference and thus shapes the reaction to deviations from that reference. Movers from expensive to cheaper cities, who have been previously exposed to high prices, spend more on housing than comparable locals (Loewenstein and Simonsohn 2006). Past exposure to generous offers in the ultimatum game causes subjects to view a given offer as too stingy and thus to reject it (Hertz and Taubinsky 2016). Priming the price of beer at resorts causes subjects to be willing to pay a higher price than if primed by the price at a shack for the very same beer to be consumed on the beach (Thaler 1985). Of course, price comparisons are essential to make good decisions, so this role of memory is helpful as well. But the use of spontaneous norms as references is sometimes misleading because the alternatives that come to mind are irrelevant, be they rents in a totally different city, payoffs in an earlier game, or locations for buying beer in determining the willingness to pay.

Where does this dual role of memory come from? Why is the automatic recall of past experiences sometimes helpful for making better choices, but sometimes counterproductive? To make progress, we propose a new model of choice based on the following idea. When we consider a good, our brain spontaneously forms a memory-based norm of its price and quality on the basis of situational cues. If the observed reality is close to the norm, we decide based on these spontaneously retrieved memories, with minimal adjustment. Equivalently, the choice problem is assigned to the "normal category" and memory acts as an anchor for valuation and choice. If instead the retrieved norm is very different from the observed reality, our attention is directed to the surprising aspects of the good. In this case, the choice does not belong to the normal category and the norm acts as a reference for adjusting our representation and valuation.

This approach to choice has several antecedents in psychology and neuroscience. Most directly, it captures the process that Kahneman and Miller (1986) describe in "Norm Theory". In their view, events spontaneously bring to mind norms by "acting as a reminder of similar events in the past". Our decisions, then, strongly react to surprise, defined as a situation in which the event is very different from the norm it evokes. Cues trigger a memory based norm, and a large discrepancy with reality triggers a large adjustment.

Similar themes appear in the psychological theories of "adaptation level" (Helson 1947, 1964), and "assimilation" vs. "contrast" (Sherif, Taub and Hovland 1958, Stapel and Sulz 2011). Adaptation level theory refers to the phenomenon that judgments are made relative to a "relevant" norm or reference, which in turn is a function of similar stimuli experienced in the past.[2] Assimilation refers to the phenomenon whereby judgments move in the direction of a memory cue or primed category. Contrast refers to the phenomenon whereby judgments move in the opposite direction of the latter. In a famous experiment (Bless and Schwarz 2010), priming subjects with "Nixon" increases the assessed corruption of a generic politician, causing assimilation to the cue. Yet the same cue "Nixon" causes subjects to judge a specific ethically problematic politician – "Newt Gingrich" – as less corrupt than if assessed without a Nixon cue, akin to a form of contrast. In our approach, these phenomena emerge as the consequences of the general mechanism described above. Because memory is associative, "Nixon" primes subjects to retrieve examples of politicians with similar characteristics. This creates assimilation, so a generic politician is judged more corrupt than absent the Nixon cue. When instead the error relative to the "Nixon norm" is made available by the direct comparison with "Gingrich", a re-evaluation of the target occurs via a contrast.

These ideas are also related to an emerging paradigm in the neuroscience of sensory perception, the so-called predictive brain model (Hawkins and Blakeslee 2004, Friston 2010, Clark 2013). In this paradigm, the brain continuously and spontaneously forms predictions of the environment. When the

[2] Helson introduced the notion of a context-specific adaptation levels into psychological judgments by expanding on his earlier work on the psychophysics of vision (Helson 1947). He also explored the determinants of the adaptation level and its implications for perception: "for every excitation-response configuration there is assumed a stimulus which represents the pooled effect of all the stimuli and to which the organism may be said to be attuned or adapted. Stimuli near this value fail to elicit any response from the organism" (Helson 1947).

prediction is good, in the sense that it explains reasonably accurately the bottom up sensory signals, it is used as the mental representation for choice. Such a representation is imperfect, but it offers a parsimonious guide to action. When instead the error is large, attention is directed to it, and the initial representation is adjusted. This paradigm accounts for a range of evidence, from the architecture of the neo-cortex to the dynamic responses of individual neurons and neuron populations to stimuli (Clark 2013, Desimone and Duncan 1995). The predictive brain model has not been evaluated for higher-level cognitive tasks such as economic choice. But the connection to sensory perception is a useful analogy that guides our analysis.

Our approach to choice seeks to bring these ideas to economics. We capture the basic intuition that valuation and choice require the recall of previous instances of similar consumption, and in particular that there is no such thing as valuation based on objective attributes independent of recalled experience. Our formal model combines memory-based norms and the error-driven allocation of attention that in turn influences choice. This approach helps account for a great deal of extremely diverse experimental and field evidence. To formalize memory-based norms, we use the textbook model of memory from Kahana (2012). Episodic memory is a database of past experiences, and external cues filter the database by retrieving past experiences that are similar to the cue while at the same time interfering with retrieval of experiences that are dissimilar from the cue. We formalize the allocation of attention and choice using the Salience theory from Bordalo, Gennaioli, and Shleifer (BGS 2012, 2013). Salience Theory holds that, when evaluating a choice option, our attention focuses on features that are very different from a reference point, or norm, while neglecting features that are similar to it. Salience can thus capture both inattention to small errors and the overweighting of "surprises" as in Norm Theory.

We develop our analysis in five steps. In Section 2, we briefly summarize the research on the psychology of memory and attention that motivates our approach, and present the central equation that describes valuation as a two-stage process of anchoring and adjustment.

In Section 3, we introduce our formal model of norms retrieved from memory, and their function as anchors for valuation. A choice provides cues that trigger recall of past experiences, and these cues can be contextual – such as buying beer at a resort – but can also comprise hedonic attributes (a specific brand or price). Because cues retrieve from memory instances that are similar to them, norms are tilted towards the cues, and the extent of adaptation to the latter depends on past experiences. When hedonic attributes are not observed, decision makers retrieve qualities and prices associated with the current context, and use these norms to assign values to their options. This analysis already sheds light on several empirical findings on assimilation to memory, including experience effects, attribution bias, projection bias, and inattention to shrouded attributes – all of which are distortions in valuation arising from selective memory.

In Section 4, we complete the model by considering the case where decision makers observe the good's quality and price, which also act as cues. Due to similarity based recall, the mind tends to retrieve accurate norms for hedonic values. In stable situations, the retrieved norms are accurate, there is no error, choice looks normal, and a rational decision is made. We then analyze two mechanisms that cause inaccurate norms and hence distorted decisions. The first stresses the role of the database of stored experiences. The second considers inaccurate norms created by particular price or quality realizations.

We find that the role of the memory database can shed light on a range of puzzling evidence on choice. One body of evidence includes background contrast effects in consumer choice (Simonson and Tversky 1992) and in ultimatum games (Herz and Taubinsky 2017), as well as Thaler's (1985) famous beer-on-the-beach experiment. This evidence illustrates the role of contrast: past experiences of high prices cause a given current price to be perceived as lower. Psychological work on assimilation and anchoring also notes the opposite phenomenon: judgments of a target stimulus move toward a prime when the prime is similar to the target (Sherif, Taub and Hovland 1958, Tversky and Kahneman 1973). Our model offers a unified account of these effects: contrast arises if memory retrieves an inaccurate norm so that attention is drawn to the surprising attributes, while assimilation prevails when the norm is similar to choice. Context plays a critical role: superficial contextual similarity, say across two stages of an experiment, may induce spuriously retrieval of the prior experiences (that are different by design), which creates artificial contrast.

In Section 5, we examine the second mechanism, namely the extent to which a good's price or quality can distort decisions by evoking an inaccurate norm. We find that frequently observed prices retrieve norms that are similar to them. The ensuing surprise is small and fails to attract attention. In this case, assimilation towards the norm creates under-reaction to small price changes. This helps explain evidence from marketing that demand is highly insensitive around prices close to the normal level (Koschate-Fischer and Wullner 2017) and evidence of inertia in consumer choice (Chandra, Handel and Schwartzstein 2018). In contrast, infrequent and extreme observed prices generate large errors, because the norm is still anchored to the moderate modal price. These errors attract attention, creating overreaction to prices in either tail. This may help shed light on why firms tend to implement frequent but small price hikes, and infrequent large discounts (Nakamura and Steinsson 2008).

Overall, similarity-based recall typically generates accurate norms and good decisions, at least to the extent that similar situations were experienced in the past. A broad experience, in the sense of increased variance of experienced prices, facilitates adaptation at any given price level. Sometimes, however, problems arise. If experienced prices lie in a narrow range, a very different price causes surprise and overreaction. While this type of error eventually fades with experience, contextual cues may be more problematic. They could be normatively irrelevant (like the sourcing of beer in Thaler's example) and yet by superficial similarity bring to mind inaccurate norms, distorting valuation. A central message of our model is that many choice anomalies may reflect superficial contextual similarity across fundamentally different experiences.

In Section 6, we connect these ideas to the literature on fairness, and in particular to some of the findings of Kahneman, Knetch, and Thaler (KKT 1986). KKT argue that a key feature of fair behavior is that it is normal, so that it easily comes to mind, and not necessarily "just" in some abstract sense. This interpretation naturally connects judgments of fairness to memory. Our model can throw light on why we see so much instability in what is perceived as fair behavior. Why, for example, do people get upset by wage cuts, but less so by bonus cuts, or why do they get upset if prices of snow-shovels increase during a snow storm but not about the prices of flowers rising on Valentine's Day? Again, part of the answer may be the fact that, despite being fundamentally similar, these situations trigger different norms because of superficial dissimilarity to past experiences. The norm for what is fair changes with context, and behavior follows.

In Section 7, we connect our approach to closely related research on memory, reference points, and expectations. Section 8 concludes with a discussion of some open problems.

2. Norms, Errors, and Choice

We begin by illustrating the structure of our model with an example, and then introduce the motivating psychological evidence. Consider a thirsty traveler arriving at an airport, and deciding whether or not to buy a bottle of water. We look at two cases: a traveler passing through security and deciding whether to go look for water (before seeing the price), and the traveler already at the airport shop and looking at the price. In the standard economic model, this is not a challenging problem. A traveler passing through security forms his best prediction of the water's quality and price $(\hat{q}, \hat{p})$ and goes to the shop if $\hat{q} \geq \hat{p}$ (assuming he has linear utility, as we do throughout this paper). If the traveler is already at the airport shop, he sees quality q and price p, and buys water if $q > p$. These two cases will be the focus of our analysis.

The standard economic model does not specify where the estimates of price and quality come from when the traveler passes through security, but presumably they come from memory. Perhaps at a deeper level, even a traveler at the shop must rely on his memory to assess whether it is worth it to him to buy a bottle of water. Of course, it matters how thirsty he is, but he must also estimate – from past experience – how much he likes the taste of bottled water, whether it will quench his thirst, but also – again from past experience – what he can do with the money if he does not buy the water. In a fundamental way, the traveler's valuation and decision about buying water depend on his past experience.

This simple example raises two questions that will be central to our analysis. First, how are past prices retrieved and how do norms or expectations form based on past experience? Second, what is the reaction to a surprise, or prediction error? To place structure on the problem, and develop predictions based on the consumer's history, we rely on the psychology of memory and attention.

2.1 Memory

Since the 1880s, a large body of experimental work has examined episodic memory, or memory of past experiences. The evidence shows that recall is a spontaneous and subconscious process in which a current experience, the cue, stimulates retrieval. As described by Kahana (2012), recall obeys two principles:

1. It is associative, driven by *similarity*: presenting a stimulus facilitates recall of items from memory that are similar to that stimulus.
2. Recall is subject to *interference*: recall of items of given similarity to the stimulus is weakened or blocked entirely in the presence of more similar items in working memory.

These principles are illustrated by two main experimental paradigms. In *item recognition* tests, subjects assess whether given words were part of a previously shown list. These tests illustrate the role of similarity because i) the probability of recall is higher for items that actually belong on the list (that is, when the cue is more similar to the item), and ii) subjects are more likely to mistakenly recognize words that are similar to a list member (they recognize yogurt when milk is on the list). In *cued recall* tests, subjects retrieve words that are pairwise associated with a cue, having previously been shown lists of relevant word pairs. These tests reveal the role of interference: if the cued word appears in many word pairs (so it is associated with many words), recall of each association is less likely. This is known as the fan effect (Anderson and Reder 1999). Interference effects are stronger for items that are more similar, so similarity shapes cued recall as well.

In Section 3 we present a model in which a consumer considering a good automatically forms a memory based prediction (q^n, p^n) through a retrieval process that follows these principles. To this end, we adapt to an economic setting the standard model of similarity based recall (Kahana 2012). In this model, past experiences are encoded as vectors of attributes. Some attributes are inherent to the event, others are contextual. For instance, when we buy water downtown, our memory records the brand and price of water, which are intrinsic or hedonic attributes, but also contextual conditions such as the time and location of the experience. Episodic memory is then described as an *event x attribute* matrix database of past experiences.[3]

In this model, one cue for retrieval is context. The feeling of thirst and the downtown location prime selective recall of past water experiences in similar locations. But intrinsic attributes such as quality or price also act as cues. Seeing the price of water at a downtown store spontaneously cues recall of past experiences, especially if they are similar. These features of recall have important implications for our theory. First, recall of similar past prices or qualities is a form of adaptation of norms towards the current good, favoring an accurate representation. Second, superficial contextual similarity may bring to mind past prices or qualities that differ from the present choice, creating an inappropriate norm.

2.2 Prediction Error, Attention and Valuation

In the simplest case in which the traveler is passing through security, he decides on whether to visit the shop based solely on the memory-based norm (q^n, p^n), so he goes to the shop if and only if $q^n \geq p^n$. But consider the case in which the consumer is already at the shop at the airport, so that the actual brand and price of water (q, p) are available to him. Is the value worth the price? To gauge the utility of drinking *q* and the disutility of paying *p*, the consumer automatically retrieves from memory past hedonic experiences of purchasing water (q^n, p^n) that are cued by the stimulus (q, p). As in the predictive brain framework, (q^n, p^n) is a spontaneous representation of the stimulus, capturing the "normal" quality and price. This representation constitutes the initial anchor for the consumption decision. If the consumer's sensory system perceives no significant difference between (q^n, p^n) and the underlying sensory stimulus

[3] We abstract from semantic memory, a broad term that covers functional associations and rule based thinking (e.g., recalling "glass" after seeing "milk"). Semantic memory allows humans to create mental models, and plays an important role in predictions.

(q, p) that generated it, the situation is "normal" and the consumer essentially chooses based on (q^n, p^n). This is a fast and frugal decision. If instead the mental representation of the choice is very different from the stimulus, we say surprise is large. In this case, the consumer's attention is driven to the error itself, leading to an adjustment of the initial representation (q^n, p^n) in the direction of the stimulus (q, p).

To give a standard example, the temperature of the doorknob of the door of our office is always slightly different, but as long as it is in the "normal range" we do not consciously perceive it. In the predictive brain framework, when our mental representation is good, top down mechanisms direct attention away from error (Hawkins and Blakeslee 2004). Relatedly, research in sensory perception has identified minimum perception thresholds below which changes in stimuli are not felt. When instead the temperature of the doorknob is very different from normal, our attention is directed to it, and we adjust our representation.

In the case of the traveler buying water, if the price at the airport is close to what he is used to, he will not pay attention to the consequences of small price differences. In fact, spontaneous representations are often good enough that choices are automatically made in the normal mode. If, however, the consumer has never seen airport prices before, his attention is directed to the large difference between the price he sees in the shop and the average price he is used to pay downtown. The consumer's attention is directed to this difference and the valuation is adjusted toward (q, p). As in Norm Theory, we strongly react to surprise.

We formalize this memory based construal of preferences by assuming that valuation is given by:

$$q^n - p^n + \sigma(q, q^n) \cdot (q - q^n) - \sigma(p, p^n) \cdot (p - p^n). \qquad (1)$$

Valuation is anchored to the norm and the errors, or surprises, are weighted by attention. The allocation of attention, and thus the weighing of errors, is modulated by the salience function $\sigma(x, x^n) \geq 0$, where x is a generic attribute. In line with our previous work (BGS 2012, 2013), the salience function $\sigma(x, x^n)$ reflects key properties of sensory perception: it increases in the proportional difference between

x and x^n. This assumption captures the idea that our attention is drawn to unusual aspects, measured in log terms, as per the well-known Weber-Fechner law.

In line with the psychology of assimilation and contrast, Equation (1) can be viewed as reflecting the use of categories to make utility evaluations. In this interpretation, the norm is the exemplar of the category of normal bottled water. If the current stimulus (q, p) falls within the normal range, it is assimilated to it, as in categorization models where perception of category members gets biased toward the average (e.g. Mullainathan 2000, Fryer and Jackson 2008). If instead the current stimulus is far from the normal category, it is placed in the "cheap" or "expensive" category, so its valuation is shaped by contrast. [4]

The remainder of the paper is organized as follows. We present the model of memory in Section 3, and show that the anchoring of valuation to norms already sheds light on a variety of puzzling evidence. In Sections 4 and 5, we examine the implications of Equation (1) for choice. This allows to unify the explanation of several puzzles that concern the assimilation vs. contrast effect of memory primes, but also findings on over and under reaction to data. Our model of recall allows to identify conditions in which different phenomena occur. In Section 6, we discuss the relationship between our model and research on fairness.

3. Memory

To present the formal model, we first describe the memory database and the process of cued recall and the formation of memory-based norms. We then show that, even without attentional mechanisms, this model can shed light on a range of well-documented memory effects in choice such as experience effects, projection bias, attribution bias, and inattention.

---

[4] This intuition or a "normal category" is also present in the predictive brain model, where the representation of the stimulus is coarser than the stimulus itself, and is compatible with a range of realizations. Similarly, in Norm Theory the norm is a range of possible values with weights that reflect their similarity to the cue. Our specification of a single point norm and of the salience function to measure distance from the norm captures this intuition and greatly improves tractability.

Episodic memory is a database of past choice experiences. Experiencing a quality-price option (q,p) in context c creates a trace $e=(q,p,c)$ in the consumer's database. Components q and p identify the hedonic attributes of an option, c captures non-hedonic attributes present during encoding, such as location or time. The experience is broader than purchasing. Considering the good in a shop, seeing its price in an advertising campaign, or being told by a friend about it, all leave traces that are potentially available for recall, though we do not allow for rehearsal of past options (Mullaintathan 2002). We assume context c to be ordinal and restrict it to the most relevant dimension for the application at hand. In reality, c is multi-dimensional and categorical. The model can be extended to capture this richness. [5]

The memory database at time t is summarized by a good-specific distribution $F_t(q,p,c)$ measuring the frequency with which past experiences of this good entailed a quality below q, a price below p, and a context below c. As new experiences come in, the distribution is updated. When context c captures calendar time, the set of times stored in memory expands. To simplify, we focus on a stable dimension c such as location, and on frequently repeated situations in which $F_t(q,p,c)$ has converged to an invariant distribution $F(q,p,c)$. We use the shorthand $F(e)$ for $F(q,p,c)$.

Because the database is good-specific, we rule out some associations. For instance, a bottled water may prime a substitute category such as soft drinks. A given price may bring to mind other goods one can purchase with the money. Our model can capture these phenomena, because the memory database and recall may be defined over the universe of goods and experiences, but we do not deal with this here. We also abstract from the possibility that surprising events leave stronger traces and thus are more easily retrieved, as in the "peak-end" rule of recall (Kahneman et al. 1993).[6]

3.1 Cues, Similarity, and Norms

[5] Whether decision makers process and remember prices is an important topic in the marketing literature. Dickson and Sawyer (1990) survey shoppers in a supermarket about their knowledge of the prices of the goods in their baskets. In their survey, 21% of shoppers do not recall the price but 56% state a price that is within 5% of the correct price.

[6] Although easier to recall, surprising events are hardly a driving force of norms. A meal at an extraordinary restaurant is memorable, but it does not alter our norm for restaurant meals, which is based on recall of more ordinary meals.

Choice at time t starts with a cue κ_t, which identifies the good (e.g., "evaluating the utility of a bottle of water") and thus the database $F(e)$ of relevant experiences, along with additional contextual detail. Often, the cue is the full current experience of observing a bottle of water q_t at a price p_t in a location c_t, which corresponds to the notation $\kappa_t = e_t = (q_t, p_t, c_t)$. This case captures the traveler who sees quality and price at the airport shop. Other times, only context is observed, so that $\kappa_t = c_t$. This is the traveler thinking about water while passing through security, where only the "airport" cue is available.

A generic cue κ_t stimulates recall of similar past experiences, where similarity acts on the set of attributes defining the cue. The cue is written $\kappa_t \equiv (\lambda_{q,t} q_t, \lambda_{p,t} p_t, c_t)$ where the indicator $\lambda_{x,t} = 1$ if attribute x is observed at time t, and is zero otherwise (it is natural to assume that context is always observed, including the category of the good under consideration). Following the standard assumption, we define similarity in terms of the geometric distance in attribute space.

Definition 1 *The similarity of a past experience* $e \equiv (q, p, c)$ *to the cue* $\kappa_t \equiv (\lambda_{q,t} q_t, \lambda_{p,t} p_t, c_t)$ *is given by*

$$S(e, \kappa_t) \equiv S(\lambda_{q,t}|q_t - q|, \lambda_{p,t}|p_t - p|, |c_t - c|) \tag{2}$$

where $S: \mathbb{R}_+^3 \to \mathbb{R}_+$ *decreases in each of its arguments.*

A context cue $\kappa_t = c_t$ retrieves past experiences based only on contextual similarity.[7] A full cue $\kappa_t = e_t$ recruits past experiences based on similarity along all attributes. Following Kahana (2012), we later use the exponential and symmetric specification:

$$S(e, \kappa_t) = \exp\{-\delta[\lambda_{q,t}(q_t - q)^2 + \lambda_{p,t}(p_t - p)^2 + (c_t - c)^2]\}, \tag{3}$$

where $\delta \geq 0$ captures the importance of similarity in recall.

The cue κ_t activates past experiences to different degrees, depending on similarity. Formally, we view this process as a cue-driven change of measure in the historical distribution $F(e)$.

Definition 2. *Denote the relative similarity of experience* e *to the cue* κ_t *by:*

---

[7] Equation (1) follows multidimensional scaling (Torgerson 1958) in which the weights capture the unequal salience of different attributes. Tversky (1977) highlights cases in which similarity does not follow geometric properties.

$$\tilde{S}(e,\kappa_t) = \frac{S(e,\kappa_t)}{\int S(e,\kappa_t)dF(e)}. \tag{4}$$

The quality and price norms for cue κ_t are the quality and price components of the average experience weighted by relative similarity to the cue:

$$e^n(\kappa_t) \equiv \int e\,\tilde{S}(e,\kappa_t)dF(e). \tag{5}$$

As in Kahneman and Miller (1986), the norm $e^n(\kappa_t)$ aggregates past experiences filtered according to similarity with the cue. The norm satisfies two properties. First, it weighs more similar experiences more heavily. Second, the weight attached to an experience decreases in the similarity of other experiences with the cue κ_t, because $\tilde{S}(e,\kappa_t)$ denotes relative similarity. This captures interference, whereby more similar memories block less similar ones (Kahana 2012).[8]

According to Equation (5), the norm is tilted toward experiences most similar to the cue. To illustrate, we focus on the simplest case where quality is fixed, so that the norm for quality is correct, $q^n = q$. We also assume that similarity has a multiplicative structure, so that when price is part of the cue we have $S(|p-p_t|,|c-c_t|) = S(|p-p_t|)S(|c-c_t|)$, which is satisfied by Equation (2). We can then show:

Proposition 1. *Denote by $p^n(c_t)$ and $p^n(p_t,c_t)$ the price norm when the cue is context or context and price, respectively. With multiplicative similarity, if price and context are uncorrelated in $F(e)$, the observed context is irrelevant for norms. In particular, denoting by $\overline{p}$ the average price in the database we have:*

i) When the only cue is context, it retrieves the unconditional average experienced price, $p^n(c_t) = \overline{p}$.

ii) When the cue is context and price, the norm is $p^n(p_t,c_t) = p^n(p_t)$, where $p^n(p_t)$ is the price norm prevailing in the hypothetical case in which the cue is only price. If the historical price distribution $dF(p)$ is normal with variance π^2, the measure $\tilde{S}(p,p_t)dF(p)$ is also normal, with variance $\pi^2/(1+2\delta\pi^2)$ and mean:

[8] Equation (3) yields the well documented laws of recency and repetition. The "contextual drift" hypothesis states that c_t moves slowly over time (e.g., our state of mind changes slowly) so context cues recent experiences. In turn, the law of repetition follows from the fact that the distorted measure $\tilde{S}(e,\kappa_t)\cdot dF(e)$ attaches a larger weight to experiences with higher frequency $dF(e)$, which thus influence the norm more than less frequent ones do.

$$p^n(p_t) = \frac{\overline{p} + 2\delta\pi^2 p_t}{1 + 2\delta\pi^2}. \tag{6}$$

When context is uncorrelated with price, it has no effect on price norms. Consider a consumer purchasing water downtown. His experiences accumulate in contexts that differ in a myriad of attributes such as brand, location, layout, etc. This bewildering variability across downtown locations, however, does not influence his price norm, because it is unrelated to the price of water. When cued with a specific store the consumer recalls past experiences at similar stores but those entail the same average price as would be retrieved in another location. In this case, norms coincide with rational expectations.

Suppose that the consumer sees both context c_t and price p_t. Because context is unrelated to price, it is again irrelevant. Observing price is however important: it triggers recall of similar past prices. The norm still depends on the modal price $\overline{p}$, which is experienced most frequently, but due to similarity is also adjusts toward the current price p_t. Similarity is thus a mechanism for "on the fly" adaptation to the current context.

When the price is observed, the norm entails an error $p_t - p^n(p_t)$. Mechanically, the error is small if recall is strongly shaped by similarity. As $\delta \to +\infty$ the price norm adapts perfectly to any price level $p^n(p_t) = p_t$, provided the latter is in the database. Interestingly, the error is small also if historical price variation π^2 is large. When the database is very dispersed, our mind can easily make sense of any current price cue by drawing on similar past realizations. In turn, these price memories strongly interfere with recall of different prices, causing the price norm to be close (i.e., to adapt to) the observed price. When instead the database is concentrated around $\overline{p}$, any price cue retrieves a norm that is close to $\overline{p}$ and entails a possibly large error or surprise.

To see this, consider again the traveler at the airport. If his memory database consists only of the deterministic downtown price of water $\overline{p}$, so that $\pi^2 = 0$, then seeing the high airport price p_A can only trigger the recall of the downtown price $p^n(p_t) = \overline{p}$. In this case, the surprise $p_A - \overline{p}$ is large. However, as airport experiences accumulate, p_A becomes more frequent in the consumer's database. This can be viewed as an increase in the variance of experienced prices π^2. With a broader database, the high price p_A

cues recall of similar past high prices, which facilitates adaptation of the norm towards the price cue, reducing the surprise. This mechanism can account for Kahneman's observation that "we are surprised only once". Each situation, even if ex-ante quite unlikely, pushes memories of similar events to the forefront of consciousness, facilitating adaptation and reducing surprise.[9]

In the example of the traveler, context also contributes to reducing error or surprise. At the airport, it is not just the high price of water that cues recall of high prices, but the airport location itself which is associated with high prices, encouraging their recall. When price and context are correlated we find:

Proposition 2 *Suppose that $F(e)$ is normal with mean $(\overline{p}, \overline{c})$, variances of price and context $\pi^2$and γ^2, and correlation ρ between them. Then, using the similarity function in Equation (1) the price norms are given by:*

$$p^n(c_t) = \frac{\overline{p} + 2\delta\gamma^2 \mathbb{E}_F(p|c_t)}{1 + 2\delta\gamma^2}, \tag{7}$$

$$p^n(c_t, p_t) = \frac{\overline{p} + 2\delta\gamma^2 \mathbb{E}_F(p|c_t) + [2\delta\pi^2 + 4\delta^2\gamma^2\pi^2(1-\rho^2)]p_t}{1 + 2\delta\gamma^2 + [2\delta\pi^2 + 4\delta^2\gamma^2\pi^2(1-\rho^2)]}. \tag{8}$$

The norms in Proposition 2 are adjusted to take the correlation ρ between price and context into account. This occurs through the term $\mathbb{E}_F(p|c_t)$, which denotes the average price observed at context c_t in the database $F(e)$. Contextual similarity creates a form of statistical conditioning. After experiencing an association of a context with high prices, this context primes a high price norm. As the traveler enters the airport, he recalls past airport experiences and the high price associated with them, because in the airport context $\mathbb{E}_F(p|c_t)$ is higher (Equation 7). When the price is also a cue, it interacts with context (Equation 8). As the traveler sees expensive water at the airport, he retrieves both past airport experiences and instances of expensive water. However, the adaptation toward the "rational expectations" benchmark $\mathbb{E}_F(p|c_t)$ is incomplete for two reasons. First, norms are still anchored to the modal experience $\overline{p}$. Second, the current price p_t is another cue anchoring price norms.

[9] Kahneman (2003) offers an auto-biographical example of this point. Having once seen a burning car on the side of a road, he half-expected to see it again when driving by the same spot (and would not have been surprised if he did).

Critically, context can either reduce or enhance the error $p_t - p^n(\kappa_t)$. By Equation (7), if the correlation between price and context is strong, the context based expectation $\mathbb{E}_F(p|c_t)$ is close to p_t and error is low, as when encountering a high price at the airport. But context can also raise the error, when it is imperfectly correlated with price. Suppose that by chance the consumer encounters a low price at the airport. Then the airport context hinders adaptation to this anomalous price. Recency effects create large errors in a similar way: if calendar time is a dimension of context, recent prices exert a stronger influence on norms. As a result, the traveler's price norm at the airport may be lowered by a low price seen downtown a day before, creating a systematic error.[10]

But there is a second and more pervasive way in which context can distort norms. Contextual cues can be normatively irrelevant – such as telling the beachgoer his beer comes from a resort – and yet trigger retrieval of associated prices, shaping norms and choice. Such context effects can be triggered by many features (location, mode of presentation, explicit allusions to other contexts, etc) provided those features correlate with intrinsic attributes of the good. Context is a double-edged sword. Our model of memory offers a way to shed light on such effects through the structure of similarity, showing that many puzzling behaviors in the field and in experiments can be understood as a by-product of this mechanism.

3.2 Implications

The idea that memory based norms act as anchors for valuation can shed light on several pieces of evidence on consumer behavior. Take the simplest case in which only context is observed, $\kappa_t = c_t$. This case could correspond to a valuation of costs/benefits that accrue in the future (e.g. assessing an investment). Here valuation is fully pinned down by the norm and utility is $q^n(c_t) - p^n(c_t)$. Suppose for simplicity that the consumer evaluates unobserved quality. From Equation (7) this assessment is given by:

$$q_i^n(c_t) = \frac{\overline{q}_i + 2\delta\gamma_i^2 \mathbb{E}_{F,i}(q|c_t)}{1 + 2\delta\gamma_i^2}, \tag{9}$$

[10] This can be formalized using the contextual drift model (Kahana 2012) by assuming context at t is the combination $\alpha c_t + (1-\alpha)c_{t-1}$ with $\alpha < 1$. Mullainathan (2002) is an early discussion of recency effects in economic decisions.

where the index i captures the fact the database may vary across individuals. In Equation (9) valuation based on memory increases in: i) the average quality $\overline{q}_i$ of the individual specific history, and ii) the average quality $\mathbb{E}_{F,i}(q|c_t)$ associated with the current context. We explore the implications of each factor in turn.

Distorted Database: Experience Effects, Attribution Bias

Malmendier and Nagel (2011) show that individuals who have experienced low stock market returns are less willing to take financial risk, and report worse returns expectations, than individuals who have experienced higher stock market returns (see also Malmendier and Nagel 2016). This "Experience effect" runs counter to the idea that individuals form rational expectations of future returns using all publicly available data. A related phenomenon is "attribution bias", whereby the valuation of a good to be consumed in the future is unduly influenced by specific past experiences with it. Haggag et al. (2018) show that consumers expect higher utility from going to an amusement park if during a past visit to the park the weather was good. Prompting consumers to think about weather moderates this behavior, suggesting that superficial neglect of concurrent conditions may create departures from rationality.

These phenomena are clearly related to memory. When thinking about future stock returns, investors recall the returns they experienced. An investor who experienced lower past returns will then have a lower $\overline{q}_i$ and a lower valuation in Equation (9).[11] A similar phenomenon occurs for the individual who had, due to nice weather, a good time at an amusing park. The structure of recall can shed further light on these effects. Providing subjects with cues c_t associated with past stock market crashes should exert a larger negative effect for investors who have experienced such crashes. For them, conditioning $\mathbb{E}_{F,i}(q|c_t)$ in the database is more pronounced due to a stronger correlation ρ. Likewise, priming all

[11] Malmendier and Nagel (2011) show that recent returns are weighed more heavily in decisions. This can be viewed as a context effect distorting $\mathbb{E}_F(q|c_t)$ towards returns experienced in temporally proximate contexts. This may also help account for the extrapolative expectations of stock returns (Greenwood and Shleifer 2014), as well as the distinction between the roles of information versus experience in decisions under risk taken in the lab (Hertwig and Erev 2009).

subjects with the same weather condition c_t should reduce attribution bias by interfering with selective retrieval of different, individual-specific experiences, thus leading to more homogeneous valuations.

Distorted Retrieval: Projection Bias and "Back of Mind" Inattention

Another set of puzzles, also arising when the relevant attributes are not observed, can be viewed as the product of selective retrieval rather than encoding. A prominent example is the "projection bias". Conlin et al (2007) show that catalogue orders of cold-weather items spike in very cold days, and items ordered in such days are more likely to be returned. Busse et al. (2015) show that consumers are more likely to buy a convertible car if the weather is better on the day they test-drive it. Using data from a large Chinese insurance provider, Chang et al. (2016) find that on days in which air pollution is high there is a spike in the purchase of health insurance, but that consumers subsequently cancel their contracts if during the lapse period air quality improves. Consumers appear to unduly weigh current conditions when estimating future benefits.

Equation (9) can account for this behavior as follows: upon seeing a high context c_t (say, good weather), memory automatically retrieves a pleasurable prediction $\mathbb{E}_{F,i}(q|c_t)$ of driving a convertible, which raises valuation. When weather turns bad, the consumer retrieves negative thoughts, which reduce $\mathbb{E}_F(q|c_t)$ and lead him to regret the prior choice. Loewenstein et al. (2003) model projection bias as a perception of excess persistence in valuation. In our account, this phenomenon arises as a special case of a broader memory-based anchoring (in this respect similar to assimilation in the Nixon experiment). Importantly, projection bias, attribution bias, and experience effects are not separate phenomena. They just hinge on different aspects of memory, a biased memory database and a selective retrieval process.[12]

A similar mechanism provides another view on inattention to shrouded product attributes (Gabaix and Laibson 2004), in particular those involving repeated neglect, which may reflect memory limitations. Consider the repeated neglect of add on fees or taxes. Chetty et al. (2009) document that when estimating

[12] A related phenomenon is the anchoring heuristic (KT 1973), whereby judgments of unfamiliar quantities, such as the length of the river Amazon, can be tilted towards irrelevant anchors. As an element of context, an anchor may act as a cue that retrieves target instances of similar magnitude. This may also help explain why an anchor whose magnitude is unreasonable for the target question does not increase the effect on judgment (Mussweiler and Strack 2001).

the cost of products, many consumers forget about sales taxes even though they regularly pay these taxes at the counter. One potential explanation is contextual interference: the base price is seen in the supermarket aisle, while the tax is experienced at the counter. The contextual dissociation between the two makes it harder to recall the tax when primed with the base price. Similarly, payment methods that decouple buying and paying may reduce accurate recall of prices at the moment of decisions (Soman and Gourville, 2001, Chatterjee and Rose 2011). To the extent that different price components are associated with different contexts, attention to those components can perhaps be increased by cueing consumers with the appropriate contexts. Exploiting the structure of memory and recall may offer a promising avenue to designing more effective reminders.

4. Error, Attention and Choice

We now turn to the case in which the decision maker observes hedonic attributes, $\kappa_t = (q_t, p_t, c_t)$. He spontaneously recalls price and quality norms $p^n(\kappa_t)$ and $q^n(\kappa_t)$, and the associated errors shape attention, triggering an adjustment of valuation. Equation (1) can then be restated more precisely as:

Definition 3. *Given cue* $\kappa_t = \left(\lambda_{q,t} q_t, \lambda_{p,t} p_t, c_t\right)$*, the utility from good* (q_t, p_t) *is given by:*

$$q^n(\kappa_t) - p^n(\kappa_t) + \lambda_{q,t}\sigma\left(q_t, q^n(\kappa_t)\right) \cdot [q_t - q^n(\kappa_t)] - \lambda_{p,t}\sigma\left(p_t, p^n(\kappa_t)\right) \cdot [p_t - p^n(\kappa_t)]. \qquad (10)$$

To formalize the predicates of Weber-Fechner law of perception, we follow our previous work and assume that the salience function $\sigma(x, y) \geq 0$ is symmetric, homogeneous of degree zero, increasing in x/y for $x \geq y \geq 0$, and that $\sigma(y, y) = 0$. We also assume that salience can be large but is bounded, in the sense that $\lim_{x/y \to +\infty} \sigma(x/y, 1) = \sigma > 1$.

Equation (10) introduces two significant modifications in the way salience affects valuation relative to our prior work. First, salience weights are attached to deviations from the norm, not to attribute levels themselves as in BGS (2013). This captures the intuition of Norm Theory and the predictive brain hypothesis, but preserves the key properties of the original salience model.[13] Second, in Equation (10) separate attention is attached to price and to quality, while in BGS (2013) what matters is the relative salience of either attribute. Once again, this modification captures the key prediction-error intuition, so that little attention is devoted to price is it is close to the norm, regardless of attention to quality.

We examine the implications of the error from two perspectives. First, we view Equation (10) as a map describing how the valuation (q_t, p_t) depends on the norm $\left(q^n(\kappa_t), p^n(\kappa_t)\right)$, and investigate the comparative statics with respect to the norm. Section 4.1 shows that this approach can account for the roles of assimilation vs. contrast. Section 4.2 uses this logic to shed light on several anomalies in consumer behavior that exploit the systematic effect of past experience on current choices. Second, Equation (10) can be viewed as describing valuation as a function of the true attributes (q_t, p_t) once the endogeneity of the norms $\left(q^n(\kappa_t), p^n(\kappa_t)\right)$ is accounted for. Section 5 then investigates the comparative statics of how valuation is influenced by the cue (q_t, p_t). This perspective sheds light on consumer under- and over-reaction to changes in the quality and price, connecting these phenomena to the adaptation of memory norms.

4.1 Assimilation vs. Contrast

When assessing a target item in light of a reference, two opposite effects are possible: the target may be assimilated to the reference, or it may be contrasted to it. In line with assimilation, the same person may be judged more competent when compared to a competent other (e.g., Pelham and Wachsmuth 1995). In line with contrast, the same car may be judged cheaper in comparison to a very expensive car (Herr 1989). Assimilation and contrast effects are pervasive, not only in judgments but also in sensory perception (Mussweiler 2003, Stapel and Suls 2011). We now show how our model generates both assimilation to, and contrast from, memory-based norms, and examine when each may occur.

[13] One advantage of the current formulation is that valuation is monotonic in price, so that a salient low price always increases valuation. In BGS (2013), by contrast, price salience reduces valuation of all goods, but reduces it more for more expensive goods. We are not aware of evidence favoring one prediction or the other. They are not critical for Salience Theory, but monotonicity creates an intuitive symmetry between the effects of high quality and low price.

For concreteness, we restrict our attention to valuation of prices. Define the valuation of price p_t in terms of norm p^n as:

$$V(p_t, p^n) \equiv p^n + \sigma(p_t, p^n) \cdot (p_t - p^n). \tag{11}$$

We call assimilation the case in which the valuation of p_t increases in the norm p^n, while we call contrast the case in which the valuation of p_t decreases in p^n. Formally, assimilation arises when:

$$\frac{\partial V(p_t, p^n)}{\partial p^n} = 1 + \frac{\partial \sigma(p_t, p^n)}{\partial p^n} \cdot (p_t - p^n) - \sigma(p_t, p^n) > 0, \tag{12}$$

while contrast arises when the reverse holds. When the error is small, $p^n \approx p_t$, the right hand side of Equation (12) is positive and assimilation prevails. In this case, little attention is paid to the error and judgment is anchored to the mental representation p^n. When instead the error is large, it attracts attention and adjustment is large. Equation (12) turns negative, reflecting a contrast with the norm.

As we show in the proof of Proposition 3, when the salience function is not too concave in the sense that $-\frac{\sigma''(x)}{\sigma'(x)} < \frac{2}{x-1}$ for $x \geq 1$ (which we assume throughout), price valuation $V(p_t, p^n)$ has the shape described in Figure 1.

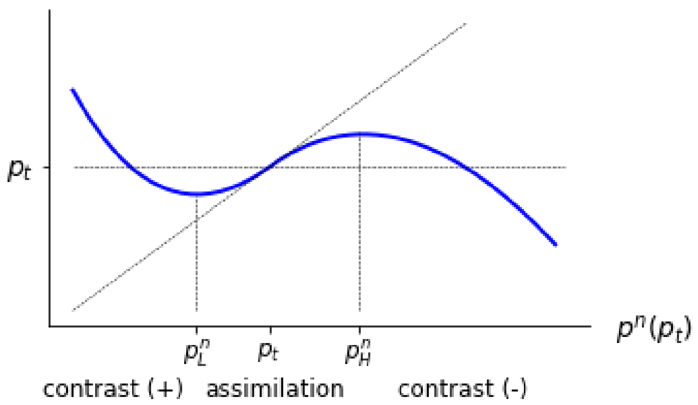

Figure 1

Assimilation, a positive relationship between the norm and valuation, prevails when the norm p^n is sufficiently close to p_t, specifically when it lies in between the two thresholds p^n_L and p^n_H. Contrast prevails when the norm is extreme. In the contrast (+) region the price norm is so low that price is exaggerated. In the contrast (-) region the norm is so high that the price is undervalued.

The range of assimilation is determined by the salience function, which pins down two thresholds for the ratio between the actual price p_t and the norm p^n that set the boundaries of the assimilation region. Factors that ceteris paribus induce the consumer to pay greater attention to the unique features of the current good should narrow down the assimilation region and make contrast more likely.

Extensive evidence in psychology confirms that contrast occurs with extreme references. Starting with Sherif, Taub and Hovard (1958) several experiments show that when judging a stimulus (e.g. the weight of an object), assimilation effects prevail when the comparison stimulus is moderate while contrast occurs when the comparison is extreme. Most of these experiments document assimilation and contrast for dimensions other than price, such as size, brightness, weight or hue (Brigell and Uhlarick 1979, Herr, Sherman and Fazio 1983). Herr (1989) shows related evidence of assimilation in the context of price: in a scale from "inexpensive" to "very expensive", subjects assimilate the price of a middle-range car to a prime of moderately expensive cars, but contrast it to a prime of very cheap or very expensive cars.[14] The conventional explanation for these effects is that assimilation occurs when the target of judgment is perceived to belong to the same category of the primed comparison stimulus while contrast arises because of sharp categorical differentiation between them. Our model offers a foundation for this mechanism based on the error entailed by the mental representation that the stimulus evokes.

4.2 Anomalies in Consumer Behavior

These very same forces can account for the effects of prior price exposure and of contextual cues on consumer behavior. Starting with Simonson and Tversky's (1992) "background contrast" effect, many experiments show that past experiences can generate reference effects. In these experiments subjects are

[14] In experiments, assimilation to a prime is more likely to prevail when the stimulus is itself ambiguous, such as when the car whose price is judged is of an unknown brand. This can easily be understood in our model: the experimental priming treatment shapes the norm only when previous knowledge of the stimulus does not interfere with it.

first presented the choice among goods characterized by certain prices. In a second stage, they choose between the same goods, but with different prices. Roughly speaking, the finding of these experiments is that prior exposure to low prices renders the current product less attractive (the contrast (+) region), whereas prior exposure to high prices makes it more attractive (the contrast (-) region).

To see how this effect arises in our model, suppose that the initial training stage endows the consumer's database with a normal price distribution $F(p)$ characterized by average price $\overline{p}$. For simplicity we momentarily abstract from context (which is akin to assuming no correlation, $\rho = 0$). In the second stage of the experiment, the consumer chooses between a good of standard quality q at price p_t and buying nothing $(0,0)$. The consumer purchases the good if and only if:

$$q - p^n(p_t) - \sigma\big(p_t, p^n(p_t)\big) \cdot \big(p_t - p^n(p_t)\big) \geq 0, \tag{13}$$

where the price norm is retrieved from memory on the basis of similarity. To see the role of prior exposure, we compute the range of prices for which Equation (13) holds. We then assess the conditions under which high past prices $\overline{p}$ boost willingness to buy as in contrast effects. Because valuation decreases in p_t, this means finding the consumer's willingness to pay, the price p_{WTP} at which Equation (13) is zero.

Proposition 3 *There are two threshold values p_* and p^* for the experienced price $\overline{p}$, defined by equations $p^n(p_L^n) = p_L^n$ and $p^n(p_H^n) = p_H^n$, such that the consumer's willingness to pay increases with experienced prices, $\frac{\partial p_{WTP}}{\partial \overline{p}} > 0$, if and only if past prices are sufficiently extreme, namely $\overline{p} \notin (p_*, p^*)$.*

Previously experienced prices $\overline{p}$ influence consumer behavior in different ways, depending on how extreme they are. When $\overline{p}$ is close to the rational willingness to pay, at the marginal choice the memory based norm is fairly accurate. As a result, assimilation effects prevail. As in the Herr (1989) experiment, a consumer primed with a slightly higher price judges a current moderate price as more expensive. This reduces willingness to buy, which is reflected in the fact that willingness to pay drops for $\overline{p} \in (p_*, p^*)$.

Matters change when the historical price $\overline{p}$ is high, or anyhow extreme, relative to the rational willingness to pay. That is, when $\overline{p} \notin (p_*, p^*)$. Now retrieval of past prices entails a high price norm and a large surprise $|p^n - p_t|$. The current price looks surprisingly cheap and the consumer is more willing to buy, reflected in the fact that p_{WTP} increases for $\overline{p} \notin (p_*, p^*)$. Now contrast prevails.

The definition of thresholds p_* and p^* shows that willingness to pay displays contrast if and only if historical prices $\overline{p}$ cause the memory based norm to fall in the range $p^n \notin (p_L^n, p_H^n)$ in Figure 1. In this sense, our model shows that standard contrast effects and the background contrast effect are due to the same mechanism: priming of an extreme mental representation so that the decision marker's attention is focused on how different, rather than similar, the target object is from it. This suggests that, like the standard contrast effect, background contrast effects are not universal: they should only arise when price differences are large.

These results offer a unified explanation of several choice anomalies. Loewenstein and Simonsohn (2006) present evidence that the home rental decisions of movers to a new city are influenced, controlling for household characteristics, by the price of housing in the origin city. When moving to a cheaper city, say from San Francisco to Pittsburgh, consumers spend more on housing than comparable locals, but they spend less than comparable locals when moving to a more expensive city, say Atlanta to New York. In subsequent renting decisions, however, the movers converge to locals: they switch to cheaper houses in Pittsburgh and to more expensive ones in New York. This behavior is puzzling for the neoclassical model, but also for models of reference dependence where the reference is the rational expectation of price, since in these models the reference point should immediately adjust. In our model, in contrast, extreme rental prices experienced in the past drive contrast effects in assessing current prices. Over time, as the mover experiences the new city, his memory database changes: $\overline{p}$ converges to local prices and the contrast effect disappears. Bordalo, Gennaioli, and Shleifer (2019) replicate the Loewenstein and Simonsohn (2006) evidence using 20 additional years of data, but also test and confirm several additional predictions of the model.

Hertz and Taubinsky (2017) document related effects in the ultimatum game. In the first stage of the experiment, they randomize subjects across treatments in which they receive high versus low offers.[15] In the second stage, all subjects play standard (one-on-one) ultimatum games. Hertz and Taubinsky find that receivers exposed to the high offers in the first stage reject more generous offers in the second stage, compared to subjects who were exposed to low first stage offers. This finding can be explained by the reaction of subjects to surprise relative to a memory-based (fairness) norm.[16] We revisit the link between memory-based norms and fairness in Section 6.

These effects arise due to another important reason: the stability of context. In housing choices, past prices come to mind because the context of a house hunt is superficially similar in San Francisco and Pittsburgh.[17] In the Hertz and Taubinsky experiment, the first and second stages of the ultimatum game are similar in many respects, which causes prior offers to come to mind. This transference of normatively irrelevant experiences would be arguably less dramatic if the two stages were less similar, just as low water prices are no longer recalled at the airport as the traveler adapts to that setting. More generally, context likely plays a critical role in reference effects and may itself be a driver of assimilation and contrast. Indeed, suppose that price and context are correlated in database $F(p,c)$ as in Proposition 2. We then have:

Proposition 4 *Suppose that context and price are positively correlated, $\rho > 0$. Then, there are thresholds c^* and c_* of realized context c_t, defined by equations $p^n(c_*, p_L^n) = p_L^n$ and $p^n(c^*, p_H^n) = p_H^n$, such that the consumer's willingness to pay increases with current context, $\frac{\partial p_{WTP}}{\partial c_t} > 0$, if and only if $c_t \notin (c_*, c^*)$.*

---

[15] In the ultimatum game, a proposer is given an endowment e and offers an allocation $(e-x, x)$ to a receiver, who chooses whether to accept or reject it. If the offer is rejected, both players get 0. The literature robustly finds that small offers ($x < 0.3e$) are often rejected and that, in turn, proposers tend to offer larger shares. In Herz and Taubinsky, the first (priming) stage consists in subjects playing 15 games of either a proposer-competition or a receiver-competition version of the ultimatum game. Consistent with previous work, receivers get a significantly higher share of the endowment when proposers compete (80% vs 30%).

[16] Formally, after a monetary offer x, the receiver in the ultimatum game is effectively choosing between accepting the good $(0, -x)$ and obtaining nothing (0,0). The monetary offer x is equivalent to a negative price. Equation (13) then implies that the offer is accepted provided:

$$u(0,-x) = x^n(n) + \sigma\big(x, x^n(n)\big)[x - x^n(x)] \geq 0.$$

In line with Proposition 3, subjects exposed to a very high previous offer retrieve a contrasting norm $x^n(x)$ and are more negatively surprised by any given offer x, which reduces their willingness to accept.

[17] Another example of superficial similarity is Genesove and Mayer's (2001) finding that sellers of condominiums anchor their selling price to their (nominal) purchase price.

When context and price are correlated, by Proposition 2 the price norm depends also on the currently observed context c_t. As a result, priming one context or another influences the norm and hence the size of the error and the extent of surprise. When context is intermediate, $c_t \in (c_*, c^*)$, it induces a norm that is close to the rational willingness to pay. In this case, assimilation prevails. Priming a context associated with slightly higher prices increases estimated disutility of buying, reducing willingness to pay. When instead context is extreme $c_t \notin (c_*, c^*)$, it induces an extreme price norm, creating contrast. Priming a context associated with higher prices reduces the current price by comparison, increasing the consumer's willingness to pay.

This role of context accords with considerable evidence from psychology. Mussweiler (2003) shows that assimilation is facilitated when similarities between the target and the comparison are made available, while contrast is enhanced when dissimilarities are made available. In the same vein, in our model similar contexts prime subjects to think of similar prices, facilitating assimilation, while dissimilar contexts do the opposite. In this sense, similarity precedes perception.

In the Hertz and Taubinsky experiment, and in the background context experiments in general, we can think of context c_t as being stable across the two experimental stages. The training stage then creates an association of this experimental context and a high price level. In the second stage, the same context is present and cues retrieval of the first stage prices. Because these prices are large relative to the current ones, we are outside of the intermediate region and subjects display context effects.

This mechanism of context-induced artificial surprise also underlies Thaler's (1985) famous beer-at-the-beach experiment.[18] Subjects state a higher willingness to pay for a given beer, to be consumed on the beach, when the beer is described as being bought from a nearby resort rather than from a nearby run-down shack. Naturally, the resort cue brings to mind high prices, while the shack cue brings to mind low prices. Contrary to the rational benchmark, then, in Thaler's experiment the stated willingness to pay for a given quality q depends on context and in particular increases with price norms. According to Proposition

[18] In related experimental evidence, consumers who are primed with "Walmart" spend less on a subsequent choice than consumers primed with "Nordstrom" (Chartrand et al 2008). Relative to Thaler (1985), this design removes a confound in that the purchase decision is exactly the same.

4, this contrast effect arises because the resort is an extreme cue. By cueing very high prices, it makes moderate prices of beer look low, raising the willingness to pay.

The resort cue here distorts behavior because it is normatively irrelevant, for the beer is anyhow consumed on the beach. It distorts choice because, by creating superficial similarity with a resort experience, it changes the standard by which the price of a beer is judged. This is an instance of the perverse role of context we discussed in Section 3: it creates artificial surprise, thereby distorting decisions. At the same time, context cues retain their usefulness. When seating at the resort, recall of shack prices would make us gasp, preventing us from enjoying the comfortable atmosphere and the attentive service. Here, the interference of the resort location with the recall of low prices reduces surprise, thereby improving decisions.

We conclude this section by discussing cases of contrast that require errors that are large enough to create overshooting $\sigma(x, x^n) > 1$ (our results up to now do not require this). One such example is constituted by "misleading sales", which refer to strategies whereby regular prices are highly inflated to lure consumers into buying during artificial sales events (BGS 2013). The high regular price primes a high price norm p^n. Critically, this norm compels the consumer to buy at the sales price $p_t < p^n$ only if the valuation of the discounted price p_t drops relative to the rational case. This is equivalent to:

$$V(p_t, p^n) < p_t \Leftrightarrow \sigma(p, p^n) > 1. \tag{14}$$

When the low price is so surprising, valuation overshoots as in a perceptual aftereffect.

Another dramatic instance of overshooting, this time in the quality space, is Hsee's (1996) broken dishes experiment. Subjects reported higher willingness to pay for a set of 10 intact dishes than for a set of 13 dishes with 11 intact and two broken. This behavior violates monotonicity. Our accounts for this choice by viewing the number of dishes as a contextual variable that retrieves an immediate quality norm from memory. In this quality norm, no dish is broken because we never experience buying broken dishes. Thus, the consumer is highly negatively surprised by the broken dishes. His valuation overshoots, and falls below the rational valuation of a set of 10 intact dishes, for which there is no negative surprise. [19]

In sum, this Section makes two broad points. First, the structure of Equation (1) helps account for the role of norms in creating assimilation and contrast effects in evaluating perceptual stimuli. Second, associative memory puts structure on the norms that are used to evaluate these stimuli in different circumstances. The analysis shows that even though contextual similarity aids decisions by favoring the retrieval of good representations, it can sometimes superficially activate poor representations that create artificial surprise, and distort decisions. Several choice anomalies can be viewed as coming from this effect.

We next explore the implications of our model for the sensitivity of valuation to a good's intrinsic attributes (q, p). In effect, we examine the comparative static with respect to the perceptual stimulus (q, p), which affects choice by affecting both recall and thus the norm (q^n, p^n), and the error and thus attention allocation to price and quality. This perspective sheds light on evidence about consumers' limited or excessive price sensitivity.

5. Memory and Price Sensitivity

A large literature in both economics and marketing stresses that valuation is somewhat insensitive to product attributes, even if they are visible. For instance, Chandra, Handel and Schwartzstein (2018) document that households do not react to changes in the price of their health plans. Marketing scholars have coined the term "just noticeable difference" (Monroe 1973) to capture the idea that consumers neglect small price changes. Other evidence points towards excess sensitivity to some price changes, particularly if they are large. Consider, for instance, the large demand shifts in response to retail sales (Casado and Ferrer 2013), or in response to price hikes of staple goods like gas (Hastings and Shapiro 2013).

[19] When evaluating the set with 13 dishes, $c = 13$, subjects immediately retrieve from memory the quality q_{13} of 13 intact dishes. With two broken dishes, however, the quality of the set is only equal to the value q_{11} of 11 intact dishes and valuation is given by $q_{13} - \sigma(q_{11}, q_{13}) \cdot (q_{13} - q_{11})$, which drops below q_{10} only if $\sigma(q_{11}, q_{13}) > 1$.

In economics, a conventional approach to this issue is rational inattention (Sims 2003), or the idea that paying attention to all price changes is too costly (Gabaix 2014). This approach can account for limited sensitivity to prices but has a hard time capturing over-reaction. In Section 5.1 we show that the prediction- error approach can generate both phenomena. In our model, attention depends on recall and in particular on the ability to retrieve accurate price norms. In Section 5.2 we discuss the implications of memory based norms for the domains of inattention and over-reaction.

5.1 Under-Reaction and Over-Reaction

To simplify, we again restrict our analysis to the case in which quality is deterministic, so that only the assessment of price is distorted. We then study the price valuation function $V\big(p_t, p^n(p_t, c_t)\big)$ as a function of p_t. It is easy to see that price valuation monotonically increases in p_t: $\frac{\partial V(p_t, p^n(p_t, c_t))}{\partial p_t} \geq 0$. Under-reaction or inattention then arises when valuation increases less than one for one with p_t, that is, when:

$$\frac{\partial V\big(p_t, p^n(p_t, c_t)\big)}{\partial p_t} < 1. \tag{15}$$

Over-reaction arises when the reverse inequality holds. Under the previous assumption that the salience function is not too concave, the Appendix shows that price valuation has the shape depicted below.[20]

[20] Figure 1 illustrates the case in which p^H is finite. This case arises when the salience determined constant $z \equiv \sigma^{-1}(1) > 1$ is sufficiently close to one, which requires that salience is sensitive enough to the percentage error. Proposition 3 obtains under the realistic assumption that (p_t, c_t) are such that all relevant norms are positive. The restriction is made necessary by the use of normal distribution and the fact that prices are restricted to be non-negative.

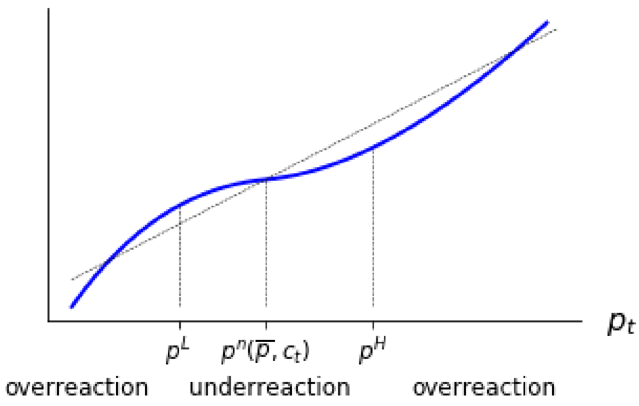

Figure 2

Valuation can be above or below the true price. The price is perceived correctly (price valuation crosses the 45-degree line) at three points. The first is the middle point p^* at which the norm is correct, identified by the equation $p^n(p^*, c_t) = p^*$. The other two points are identified by perfect attention-based adjustment of the initial norm, namely $\sigma\big(p_t, p^n(p_t)\big) = 1$.

Under-reaction prevails when the observed price is close to the prediction from memory, $p_t \in (p^L, p^H)$. In this range, the valuation function is too flat. Here the decision-maker is inattentive to deviations from the norm because the norm retrieved from memory is fairly accurate so that error is low. As a consequence, top down feedback suppresses attention and valuation remains anchored to the memory based prediction. When instead the realized price is far from the norm, the error is large and salient, the adjustment of valuation to this error is large, and valuation over-reacts.

The inattention regime in which prices change within a moderate range can help explain the rigidity of behavior observed in some economic settings, but also reconcile it with the high elasticity observed in others. A large literature in marketing examines the elasticity of demand to price changes (Cheng and Monroe 2013). Examining aggregate demand across dozens of retail product categories and different retailers, this literature suggests that demand is rigid with respect to small price changes, but shows greater elasticity for larger price changes in either direction (Pauwels et al 2007, Casado and Ferrer, 2013). We would argue that these patterns have to do with memory. Small price changes within the "normal range" are not attended to, but large changes are, perhaps too much so.

In the context of retail markets, Nakamura and Steinsson (2008) find that price increases are on average smaller than price decreases. Selective memory could help explain these effects. A sequence of small price changes gradually changes the price norm in the same direction, yielding a sequence of small and thus unnoticeable errors. But the cumulative effect of many neglected price increases can be substantial and costly to the consumer, so that retailers may take advantage of this when raising prices. This evidence cannot easily be accounted for by switching costs on the part of consumers. This mechanism may also be at play in the evidence for inertia shown by households relative to the prices of their health plans (Chandra, Handel and Schwartzstein 2018).

The valuation of price in our model, illustrated in Figure 2, is qualitatively different from the value function proposed by Prospect Theory (KT 1979). In Prospect Theory, taking the price norm as the reference, the value function is assumed to be the steepest around the norm, and the value function for price is steeper above the norm than below it to reflect loss aversion. In contrast, in our model valuation is flatter around the norm, reflecting Weber's law and inattention to small attribute changes. Moreover, diminishing sensitivity suggests that valuation is steeper below the norm rather than above it. These properties hold symmetrically for both attributes. Consumers would be more sensitive to a price cut from \$10 to \$5 than to a price increase from \$10 to \$15, just as they would be more sensitive to a drop in quality than to a commensurate increase in quality. The evidence reviewed above, in particular the prevalence of small price hikes and occasional large price cuts (Nakamura and Steinsson 2008) is consistent with this view.

5.2 Experience, Under- Reaction, and Over-Reaction

One important implication of our model is that inattention depends on the quality of norms. When the environment is such that memory makes good predictions, error is small and the consumer is inattentive. When instead memory predictions are poor, the consumer is highly attentive to price and hence over-reacts. We next illustrate this point by showing that the structure of similarity in recall helps make predictions about the regimes of assimilation versus contrast. To do so, we focus on the role of the past distribution of prices and context in shaping the region of under-reaction (p^L, p^H).

Proposition 5. *Suppose that $F(p, c)$ is joint normal. Our model yields the following comparative statics:*

1) The range of under-reaction (p^L, p^H) moves up and expands for a consumer who has experienced higher prices on average, formally $\frac{\partial p^L}{\partial \overline{p}} > 0, \frac{\partial p^H}{\partial \overline{p}} > 0$, and $\frac{\partial (p^H - p^L)}{\partial \overline{p}} > 0$.

2) Priming the consumer with a context c_t associated with higher prices has the same effect as in 1).

3) When context is uncorrelated with prices, $\rho = 0$, higher volatility of prices π^2 expands the under-reaction range (p^L, p^H) symmetrically around $\overline{p}$, formally $\frac{\partial p^L}{\partial \pi^2} < 0, \frac{\partial p^H}{\partial \pi^2} > 0$.

It is intuitive that decision makers are inattentive to prices close to the norm. Consumers routinely face tiny price increases, but are not much affected by them. According to 1), this occurs to a greater extent at a higher experienced price level $\overline{p}$. This is due to the diminishing sensitivity of salience: an error of $5 is less salient compared to a high price norm such as $100 than to a low price norm of $10. This finding is reminiscent of Dehaene's (1996) evidence of Weber's law in number perception. Several papers in marketing offer laboratory and field evidence of a similar phenomenon in the consideration of price: the threshold at which a good is perceived as too expensive grows proportionally with the average price level of that good, generating a "latitude of price acceptance" which also grows with price (Koschate-Fischer and Wullner 2017).

According to 2), price inattention is context-dependent. A cue inducing recall of similar contexts associated with low prices tends to shrink the price inattention range and to move it to the left. At a market in a developing country, consumers may bargain hard over amounts that they would not even notice when shopping in a rich country. Dissimilarity from a rich country market shuts down recall of the traveler's home experience. Once back home, the consumer again adapts to the usual prices, but cues that bring the previous experience to mind jolt him to consider the usual prices as being high. Contextual similarity brings past price norms to mind, switching inattention on and off.

Adaptation also expands the range of inattention. Exposure to higher price variability π^2 makes our mental prediction less rigid, dampening surprise. As a result, the error becomes more likely to go unnoticed. In models of rational inattention, greater ex-ante price variability should induce the agent to pay more attention to prices. Memory entails a countervailing force: higher experienced price variability makes mental representations more flexible, reducing errors. In fact, inattention arises only for errors that are small in proportional terms. In this sense, the valuation errors entailed by inattention are naturally bounded because inattention is endogenous to the (ex post) surprise. These effects could be analyzed by studying consumer attention to prices across product categories featuring differential price variability. More generally, because the same general mechanism of memory retrieval and error creates inattention as well as surprise and overreaction, it allows for testable predictions regarding the transition between different biases.

An experiment by Niedrich, Sharma and Wedell (2001) is consistent with the role of experienced price variability in driving memory adaptation and thus reduced attention to price. The authors show subjects price series drawn from different distributions, and then ask them to report the attractiveness of the product. They find that when prices are drawn from a bimodal distribution, subjects are less attracted by low prices and less disappointed by high prices relative to subjects trained with less volatile distributions. Our model can account for this fact: after adapting to the bimodal distribution, subjects view both high and low prices as normal. As a consequence, subjects are less attentive to price and judgments of attractiveness are flatter. The same idea can explain why the efficacy of "strategic sales" is limited if these sales occur too often: consumers become adapted, and hence they are not surprised when they see them. The marketing literature argues that frequent shallow sales lower consumers' "internal reference price" much more dramatically than do infrequent deep sales (Chang and Monroe 2013). This is consistent with our model: shallow sales are both more frequent and more similar to regular prices, and thus may entail assimilation. When the regular price is posted, recall of sales prices makes the current price look like a bad deal. When the product is on sale, recall of past sales prices makes the current discount look normal.

6. Norms, Fairness and Social Preferences

A large body of work challenges the assumption of individual self-interest by showing that fairness influences economic decisions in the field and the lab (see Fehr and Schmidt 2006). These studies often rely on two ingredients: a "fair" allocation, and social preferences that make departures from it personally costly. Several models of social preferences have been developed. For instance, Fehr and Schmidt (1999) assume that an agent is altruistic toward players whose payoff is below the fair allocation and envious about the others. But several other approaches have been considered (Charness and Rabin 2002).

Much less attention has been devoted to the question of where the fair allocation comes from. Often, it is assumed to be equal split of payoffs, but this is conceptually and empirically unsatisfactory. In Rabin's (1993) model of reciprocity, fair behavior of a player is defined with respect to a counterfactual action the same player may have taken. Here we follow a different approach, and connect fairness norms to memory. We already showed some examples of this logic in the context of "fair prices". We now further show how memory based norms can help explain different, at times puzzling, instances of fairness judgments and behavior. We stress that our focus here is on the determination of the fair transaction. We have nothing new to say at this point about social preferences attaching utilities to departures from such a transaction.

In a paper published in the same year as Norm Theory, Kahneman, Knetsch and Thaler (KKT 1986) equate the fair allocation with a so-called "reference transaction", described as a norm. The authors state "the reference transaction provides a basis for fairness judgments because it is normal, not necessarily because it is just." That is, the fair allocation does not necessarily follow from abstract principles. It conforms to a norm based on experience. We would say that these norms are retrieved from memory, and largely reflect custom. This further suggests that fair allocations are malleable, and adapt to change. As KKT put it, "terms of exchange that are initially seen as unfair may in time acquire the status of a reference transaction (...) at least in the sense that alternatives to it no longer come to mind." Here fairness and availability in memory nearly coincide.

To see the implications of our model for this set of issues, consider the following KKT experiment on the instability in judgments of fairness. The authors describe an employee who has worked at a small shop for six months and earns $9 an hour. There is some unemployment in the area so similar shops are hiring reliable workers for $7 an hour. Survey respondents are asked to judge the fairness of the firm's actions in two scenarios. In scenario 1, the firm cuts the employee's wage down to the market level. In scenario 2, the employee quits and the firm hires a new worker for $7. A vast majority of respondents find the first scenario unfair, while the same respondents find both the second and third scenarios to be fair. In each case, a small modification of the problem leads to drastically different fairness judgements. Why?

We view this context-specificity as stemming from recall of past experiences, similarly to the Hertz and Taubinsky's (2017) ultimatum game experiment. In that setting, experience with generous offers in the first stage subsequently led receivers to judge as unfair, and thus reject, good second stage offers that they would otherwise accept. The superficial similarity of the two stages of the experiment encouraged receivers to recall first stage offers and use them as fairness norms. As stressed in the introduction, while contextual cues can facilitate adaptation to the current environment, they can also generate superficial similarity to past experiences that are currently irrelevant.

We think that the same effect is at work in the KKT experiments. Scenario 1 induces subjects to selectively recall similar past instances of employed workers facing fluctuations in the market wage. Here the recall of downward wage rigidity causes subjects to see a wage reduction as surprisingly atypical and, hence, unfair. By looking at Scenario 1 in isolation one may think that subjects' notion of fairness is based on distributional justice, favoring an equitable distribution of the surplus. Scenario 2 shows that this is not the case. Here, letting the firm reap the entire surplus vis-à-vis a new entrant is judged to be fair. Again, we view this as being again due to memory: Scenario 2 causes subjects to recall similar past experiences of new hires. These memories make it look normal, and hence fair, for the firm to only pay the going market wage.[21]

In this account, fairness norms are unstable because they do not reflect universal welfare principles. They are derived on the fly, through memory of past experiences. One implication is that the practice of certain customs renders them self-reinforcing. To see this, consider a price p that is normally distributed with context c as in Proposition 2. p may capture the price of snow shovels and c weather (e.g., whether there is a snow storm). Alternatively, p may capture the wage of a worker and c the market wage. Or p may stand for the price of flowers and c is a day of the year, such as Valentine's Day. Equation (7), implies that the norm changes as follows:

$$\frac{\partial p^n(c_t, p_t)}{\partial c_t} \propto 2\delta\rho\pi\kappa.$$

When context changes, a price or wage change is considered fair if: i) similarity is at play $\delta > 0$, and ii) price and context have been historically associated in the same way, as captured by ρ. The role of similarity is key. When consumers see higher context c_t, they recall a *specific* situation in which the price is typically higher, $\rho > 0$, and accept a higher current prices as fair.

By this account, raising the price of snow shovels during a storm is considered unfair due to the very custom of not raising such price, $\rho = 0$. Due to this custom, deviating from the norm causes negative surprise and refusal to buy, which makes the custom itself self-reinforcing. In other settings, it is instead considered fair to raise prices with demand. It is for instance well accepted that flowers are more expensive on Valentine's Day.[22] Even here the self-reinforcing memory mechanism plays an important role. On Valentine's day, people recall high prices of flowers because of their past experience. Likewise, while firms' reluctance to cut wages (Bewley 2007) justifies the unfairness of this practice, firms routinely cut bonuses and other benefits, which is not viewed as unfair. The fact that customs are, at least in part, self-reinforcing implies that they may be actively changed. But doing so takes time. The memory database of

[21] KKT discuss a third scenario in which a worker moves to a new task in the firm, for which pay is 7$ rather than 9$. Also in this case, paying the worker the 7$ wage is viewed as fair. Similarity based recall can account for this evidence: task-similarity primes subjects to recall the 7$ wage, deeming it fair. This is akin to judging higher prices fair at resorts.

[22] On the basis of general principles, one may even argue that since snow shovels can be purchased in advance by consumers while Valentines' day flowers cannot, it would be fairer to have inflexible flower rather than shovel prices.

consumers must adapt to new practice, until the force of similarity renders the new custom acceptable.

Another instance of instability of fairness norms comes from experiments on the dictator game. The literature has robustly shown that small modifications in the presentation of the problem can have large implications for behavior even if they are payoff-neutral (List 2007, Krupka Weber 2016). In particular, Krupka and Weber (2016) run experiments on two variants of the Dictator game. In one variant, player 1 is given a total of $10 to share; in the other variant, both players are given $5, and player 1 chooses whether to take up to, or give up to, $5. The two versions are strategically equivalent. Yet in the experiments player 2 ends up with more in the second variant than in the first, and is significantly more likely to end up with $5 in the second treatment.

Krupka and Weber interpret the evidence by saying that the first variant brings to mind a sharing frame whereas the second brings to mind a taking frame, which is socially sanctioned. This interpretation feels right and resonates with memory-based norms. In our model, the context of the second variant (taking) filters in the memory database real world situations in which there is a choice of how much to take. Because taking is often discouraged or unlawful, this context mechanically evokes a low norm, inducing player 1 to be content with a payoff closer to the initial endowment. The sharing variant evokes giving, which brings to mind situations in which something, but not a lot, is given. Behavior then changes across versions because the description of the problem cues retrieval of different representations of choice.

7. Related Literature

The dual role of memory, as an anchor and as a reference for valuation, allow us to discuss several concepts in the literature, such as reference points, expectations, and representations, from the unifying perspective of memory-based norms.

Our approach captures several aspects of existing models of reference points.[23] Consider first the classic Kahneman and Tversky's (1979) "status quo" reference point. Their idea is closely related to memory-based norms in stable settings where the memory database is overwhelmingly populated with the habitual experience. For example, in his first time at the airport, the traveler's norm is the habitual low price of water he experienced downtown. An employee's salary norm is his usual salary.

In a changing environment, it is clear that reference points adapt. Koszegi and Rabin (2006) capture this phenomenon using forward-looking rational expectations. Della Vigna et al. (2017), in contrast, use mechanically adaptive, backward-looking expectations.[24] Memory-based norms have a backward-looking component, in that they reflect and filter past experiences, but are also forward-looking in the sense that the current data and context triggers similar past experiences creating selective adaptation.[25] As such, they help shed light on when decision makers are anchored to recent past experiences and when instead they are well adapted to the current choice.

Consider next the broader notion, gaining ground in the literature, that the first step of choice is the representation of choice stimuli. Recent papers study choice under limited attention or efficient coding of noisy stimuli. Gabaix (2014) builds a model in which decision makers form a sparse representation of their environment using only a subset of the data. This model admits a representation similar to Equation (1) in which a prior about attributes is adjusted using attention weights that do not have the proportional structure of salience and that in the main specification is optimally chosen ex ante. Woodford (2012) builds a model where the sensory system is noisy, so our cognitive apparatus filters the noise optimally. This

---

[23] In most theories, the reference point is triggered by the situation, not by each good in the choice set. For instance, it can consist of the decision maker's choice, past or expected (e.g., Koszegi and Rabin 2006), or by the decision maker's consumption state (Kahneman and Tversky 1979). In Appendix X we extend our model to choice among many goods, and show how it can accommodate these effects, for instance by allowing each option to disproportionately retrieve past choices (that is, the option that was actually chosen).

[24] Bell (1985) also identifies the reference price as the rational expected price (see also Gul 1991). Barberis and Huang (2001) and Barberis and Xiong (2009) take a related approach in asset pricing with respect to the expected risk free rate.

[25] To give an example, consider again buying water at the airport. During his second visit at the airport, the consumer remembers that water is more expensive there. Under rational expectations, adaptation is immediate. Memory-based adaptation might be more sluggish. Unlike under mechanically adaptive reference points, though, similarity causes memory-based norms to adapt to different settings. That is, there are two "fair" prices of water: one downtown, another at the airport. With enough experience in each context, selective adaptation causes memory-based norms to look like rationally expected prices both downtown and at the airport.

model also admits a representation in which the prior about attributes is adjusted when a choice option is observed, with a weight equal to the signal to noise ratio of the sensory system. Frydman and Jin (2018) extend this approach by examining efficient coding and show that a previously experienced distribution of prices shapes subjects' subsequent choices, and perceptions of prices.

Our model departs from this work in three ways. First, it endogenizes the initial mental prediction, the prior, on the basis of a well-documented model of associative memory. Second, attention is allocated ex post on the basis of the error, which implies that the error in the representation is contained. Third, relative to models of inattentive choice, which rationalize rigidity, our model also captures over-reaction to surprise. Over-reaction in our framework is due to a reallocation of attention to surprising data which cannot be explained using initial mental representations. In particular, models of efficient coding predict that valuation is most sensitive in the range of most frequently experienced stimuli, in contrast to our analysis in Section 4.

There has also been recent work on positive models of attention allocation in economics. Besides salience theory (BGS 20120, 2013), also Koszegi and Szeidl (2013) and Bushong, Rabin and Schwartzstein (2017) explore how attention is allocated as a response to the choice data. Relative to the salience model, they focus on the role of the range of attributes, and in particular on whether larger attribute ranges attract or dampen attention. Endogenizing memory-based norms suggests both intuitions are valid: while extreme price realizations unambiguously attract attention (as in the water at the airport example), experiencing a greater variance of prices facilitates the decision-maker's adaptation to any price, and dampens overreaction.

We conclude with a discussion of previous work on memory in economics. In case based decision theory (CBDT, Gilboa Schmeidler 1995), decision makers assess risky choices according to their similarity with past choices. The role of similarity is reminiscent of our model. There are two main differences, though. First, in CBDT similarity is characterized axiomatically, not psychologically. For example, CBDT does not allow for contextual attributes to influence recall, a crucial feature both of the standard model of memory and of our approach, and one that is critical in explaining a lot of data. Second, CBDT does not consider the allocation of attention on the basis of surprise. In our language, in CBDT assessment are anchored to memories, but there are no reference effects.[26]

Recent work takes a more psychological approach. In Mullainathan (2002), limited memory distorts Bayesian updating and forecasting of an economic variable. Similarity influences recall, but there is neither role for context nor for interference. Taubinsky (2014) studies optimal reminders in a model where memory is imperfect and mental rehearsal promotes recall. Ericson (2016) studies the interaction of forgetting and procrastination, drawing implications for the demand for reminders. Following adaptation level theory, the marketing literature discusses the reference role of past choice, but these models focus on recency effects (e.g., Cheng and Monroe 2013). Nagel and Xu (2018) model expectation formation under limited memory, where past data is forgotten, and Silveira and Woodford (2019) derive optimal recall under limited memory. These models feature recency effects but abstract from similarity and interference. A number of recent papers build on Kahneman and Tversky's (1972) representativeness heuristic to explore how selective memory shapes beliefs. In this approach, recall if selected toward features that are most diagnostic of, or similar to, a group in contrast to a comparison group (Gennaioli and Shleifer 2010, Bordalo, Coffman, Gennaioli, Shleifer 2016). While these effects could also be at play in the recall of norms, our approach, like Norm Theory, abstracts from them.

8. Conclusion

In this paper, we propose a new theory of choice based on the idea that valuation is a two stage process: an estimation of value based on cued recall, followed by an adjustment in the direction of any discrepancy between the estimated and observed attributes. We use a biologically founded, textbook model of memory (Kahana 2012) to build a model of memory-based norms, and combine it with the

[26] Other papers explore optimal storage of information (e.g., Dow 1991, Wilson 2014) or decision problems with exogenous imperfect recall (Piccione and Rubinstein 1997). Rubinstein (1998) summarizes some of this work.

salience theory of choice, which is a natural way to incorporate the notions of surprise, and over-reaction to surprise, that are critical to Kahneman and Miller's Norm Theory.

The critical feature of the model – recall through similarity – yields many predictions on what comes to mind when decision makers face a stimulus, which have been extensively tested and confirmed in memory research but which also have multiple implications for economic analysis. Due to the central role of similarity in recall, norms tend to be adapted to the specific choice at hand, creating the stability (and even rationality) of choice that is often observed. At the same time, norms can also incorporate normatively irrelevant contextual features, and through this channel lead to unstable and apparently irrational choice. The model sheds light on when adaptation or surprise arise, and provides a unified account of a range of evidence in the literature as well as making several novel and testable predictions.

Throughout this paper, we have made a number of specific modeling choices for clarity, many of which can be revisited or relaxed. There are several missing aspects in the basic model of memory, such as the importance of salient memories, the inattention to some aspects of the initial stimulus that may influence recall (Schwartzstein 2014), or even the failure of initial encoding of some experiences. In addition, with some modifications, our model can perhaps also incorporate recall of other types of information from memory, such as goals or information about future events. Similarly, the recall cued by a choice option is not necessarily restricted to instances of the same good, but can include alternative goods or alternative uses of funds, which can shape how decision-makers frame their own decisions.

Our model draws a link between memory-based norms and expectations of value, and this raises two important issues. First, in more complex tasks such as predicting the next realization of a stochastic process, decision makers do not just recall past realizations, but also use forward looking intuitions about the data generating process. To capture this aspect requires broadening the model to include semantic memory. Second, the representativeness heuristic has been shown to be an important driver of expectations in macroeconomic, financial, and social domains. Bordalo et al. (2018) use experimental methods to show that the representativeness heuristic influences probability judgments by affecting recall, but further thought is necessary to connect the similarity model of recall with those findings. While many questions remain open, it is clear that textbook models of memory offer an opportunity to complete behavioral models and to improve their empirical testability, and at a deeper level to understand how decision makers represent and make choices.

References

Anderson, John and Lynne Reder. 1999. "The Fan Effect: New Results and New Theories," *Journal of Experimental Psychology* 128 (2): 186 – 197.

Azeredo da Silveira, Rava and Michael Woodford. 2019. "Noisy Memory and Over-Reaction to News", NBER w25478.

Barberis, Nicholas and Ming Huang. 2001. "Mental Accounting, Loss Aversion, and Individual Stock Returns." *Journal of Finance* 56(4): 1247–1292.

Barberis, Nicholas and Wei Xiong. 2009. "What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation." *Journal of Finance* 64(2): 751–784.

Bell, David. 1985. "Disappointment in Decision Making Under Uncertainty." *Operations Research* 33(1): 1—27.

Bewley, Truman. 2007. "Fairness, Reciprocity, and Wage Rigidity," in *Behavioral Economics and its Applications*, Peter Diamond and Hannu Vartiainen Editors, Princeton University Press, Princeton NJ.

Bless, Herbert and Norbert Schwarz. 2010. "Mental Construal and the Emergence of Assimilation and Contrast Effects: The Inclusion/Exclusion Model," *Advances in Experimental Social Psychology* 42: 319 – 373.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. 2016. "Stereotypes." *Quarterly Journal of Economics* 131(4): 1753 -- 1794.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice under Risk." *Quarterly Journal of Economics* 127 (3): 1243 -- 1285.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy* 121(5): 803 -- 843.

Brigell, Mitchell and John Uhlarick. 1979. "The Relational Determination of Length Illusions and Length Aftereffects" *Perception* 8 (2): 187 – 197.

Bushong, Benjamin, Matthew Rabin and Joshua Schwartzstein. 2016. "A Model of Relative Thinking." Unpublished Working Paper.

Busse Meghan, Devin Pope, Jaren Pope, Jorge Silva-Risso. 2015. "The Psychological Effect of Weather on Car Purchases," Quarterly Journal of Economics 1 (1): 371 – 414.

Casado, Esteban and Juan-Carlos Ferrer. 2013. "Consumer Price Sensitivity in the Retail Industry: Latitude of Acceptance with Heterogeneous Demand," *European Journal of Operational Research* 228 (2): 418 – 426.

Chandra, Amitabh, Benjamin Handel, and Joshua Schwartzstein. 2018. "Behavioral Economics and Health-Care Markets," working paper.

Chang, Tom, Wei Huang and Yongxiang Wang. 2018. "Something in the Air: Pollution and the Demand for Health Insurance," *Review of Economic Studies* 85 (3): 1609–1634.

Charness, Gary and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics* 117 (3): 817–869.

Chartrand , Tanya, Joel Huber, Baba Shiv , and Robin Tanner. 2008. "Nonconscious Goals and Consumer Choice," *Journal of Consumer Research* 35(2): 189–201.

Chatterjee, Promothesh and Randall Rose. 2011. "Do Payment Mechanisms Change the Way Consumers Perceive Products?" *Journal of Consumer Research* 38 (6): 1129–1139.

Cheng, Lilian and Ken Monroe 2013. “”An Appraisal of Behavioral Price Research (Part 1): Price as a Physical Stimulus,” *Academy of Marketing Science Review* 3 (3): 103 – 129.

Chetty, Raj, Adam Looney, Kory Kroft. 2009. “Salience and Taxation: Theory and Evidence,” *American Economic Review* 99 (4): 1145 – 1177.

Clark, Andy. 2013. “”Whatever next? Predictive brains, situated agents, and the future of cognitive science,” *Behavioral and Brain Sciences* 36 (3): 181 – 204.

Conlin, Michael, Ted O’Donoghue, and Timothy Vogelsang. 2007. “Projection Bias in Catalogue Orders,” American Economic Review 97 (4): 1217 – 1249.

Dehaene, Stanislas. 1997. *The Number Sense.* New York: Oxford University Press.

DellaVigna, Stefano, Attila Lindner, Balazs Reizer, and Johannes Schmieder. 2017. “Reference-Dependent Job Search: Evidence from Hungary." *Quarterly Journal of Economics* 132 (4): 1969 – 2018.

Desimone, Robert and John Duncan. 1995. “Neural Mechanisms of Selective Visual Attention” *Annual Review of Neuroscience* 18: 193 – 222.

Dickson, Peter and Alan Sawyer. 1990. “The Price Knowledge and Search of Supermarket Shoppers,” *Journal of Marketing* 54 (3): 42—53.

Dow, James. 1991. "Search Decisions with Limited Memory." *Review of Economic Studies* 58(1): 1-14.

Ericson, Keith. 2016. “On the Interaction of Memory and Procrastination: Implications for Reminders,Deadlines and Empirical Estimation.” *Journal of the European Economic Association,* forthcoming.

Fehr, Ernst and Klaus Schmidt. 2006. “The Economics of Fairness, Reciprocity and Altruism – Experimental Evidence and New Theories,” *Handbook of the Economics of Giving, Altruism and Reciprocity,* Serge-Christophe Kolm, Jean Mercier Ythier Editors.

Friston, Karl. 2010. “The free-energy principle: a unified brain theory?” *Nature Reviews Neuroscience* 11: 127–138.

Frydman, Cary and Lawrence Jin. 2018. “Efficient Coding and Risky Choice,” working paper.

Fryer, Roland and Matthew Jackson. 2008. “A Categorical Model of Cognition and Biased Decision Making,” The B.E. Journal of Theoretical Economics 8 (1): article 6.

Gabaix, Xavier. 2014. "A sparsity-based model of bounded rationality." *Quarterly Journal of Economics* 129 (4): 1661-1710.

Gabaix, Xavier and David Laibson. 2004. “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics* 121 (2): 505–540.

Genesove, David and Christopher Mayer. 2001. “Loss Aversion and Seller Behavior: Evidence from the Housing Market." *Quarterly Journal of Economics* 116(4): 1233 -- 1260.

Gennaioli, Nicola and Andrei Shleifer. 2010. “What Comes to Mind." *Quarterly Journal of Economics* 125 (4): 1399 -- 1433.

Gilboa, Itzhak and David Schmeidler. 1995. “Case-Based Decision Theory." *Quarterly Journal of Economics* 110(3): 605 -- 639.

Greenwood, Robin and Andrei Shleifer. 2014. “Expectations of Returns and Expected Returns,” *Review of Financial Studies* 27 (3): 714–746.

Gul, Faruk. 1991. “A Theory of Disappointment Aversion,” *Econometrica* 59 (3): 667-686.

Haggag, Kareem, Devin Pope, Kinsey Bryant-Lees, Maarten Bos. 2018. “Attribution Bias in Consumer Choice,” *Review of Economic Studies*, forthcoming.

Hare, Todd, Jonathan Malmaud and Antonio Rangel. 2011. “Focusing Attention on the Health Aspects of Foods Changes Value Signals in vmPFC and Improves Dietary Choice” *Journal of Neuroscience* 31 (30): 11077-11087.

Hastings, Justine and Jesse Shapiro. 2013. “Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *Quarterly Journal of Economics* 128(4): 1449 -- 1498.

Hawkins, Jeff and Sandra Blakeslee. 2004. *On Intelligence.* Times Books, New York, NY.

Helson, Harry. 1947. “Adaptation-Level as Frame of Reference for Prediction of Psychophysical Data” *American Journal of Psychology* 60 (1): 1-29

Herr, Paul. 1989. “Priming Price: Prior Knowledge and Context Effects,” *Journal of Consumer Research* 16(1): 67–75.

Herr, Paul, Steven Sherman, and Russell Fazio. 1983. “On the Consequences of Priming: Assimilation and Contrast Effects,” *Journal of Experimental Social Psychology* 19 (4): 323-340

Hertwig, Ralph and Ido Erev 2009. “The Description–Experience Gap in Risky Choice,” *Trends in Cognitive Sciences* 13(12): 517-523.

Herz, Holger and Dmitry Taubinsky. 2017. “What Makes a Price Fair? An Experimental Study of Transaction Experience and Endogenous Fairness Views,” *Journal of the European Economic Association* 16 (2): 316–352.

Hsee, Christopher. 1996. “The Evaluability Hypothesis: An Explanation for Preference Reversals between Joint and Separate Evaluations of Alternatives,” *Organizational Behavior and Human Decision Processes*

67 (3): 247-257.

Kahana, Michael. 2012. *Foundation of Human Memory*. Oxford University Press, Oxford UK.

Kahneman, Daniel. 2003. “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review* 93 (5): 1449 – 1475.

Kahneman, Daniel, Jack Knetsch and Richard Thaler. “Fairness as a Constraint on Profit Seeking: Entitlements in the Market,” *American Economic Review* 76 (4): 728-741

Kahneman, Daniel and Dale Miller. 1986. “Norm Theory: Comparing Reality to its Alternatives.” *Psychological Review* 93(2): 136-153.

Kahneman, Daniel, Barbara Fredrickson, Charles Schreiber, and Donald Redelmeier. 1993. “When More Pain Is Preferred to Less: Adding a Better End. ” *Psychological Science* 4(6): 401-405.

Kahneman, Daniel and Amos Tversky. 1972. “Subjective Probability: A Judgment of Representativeness.” *Cognitive Psychology* 3 (3): 430–454.

Kahneman, Daniel and Amos Tversky. 1979. “Prospect Theory: an Analysis of Decision under Risk." *Econometrica* 47 (2): 263 -- 292.

Koschate-Fischer, Nicole and Katharina Wullner. 2017. “New Developments in Behavioral Pricing Research,” Journal of Business Economics 87 (6): 809–875.

Koszegi, Botond and Matthew Rabin. 2006. “A Model of Reference-Dependent Preferences", *Quarterly Journal of Economics*, 121 (4): 1133 -- 1165.

Koszegi, Botond and Adam Szeidl. 2012. “A Theory of Focusing in Economic Choice” *Quarterly Journal of Economics* 128 (1): 53-104.

Krupka, Erin and Roberto Weber. 2013. “Identifying Social Norms Using Coordination Games: Why Does Dictator Game Sharing Vary?” *Journal of the European Economic Association* 11 (3): 495–52.

List, John. 2007. “On the Interpretation of Giving in Dictator Games,” *Journal of Political Economy* 115(3): 482-494.

Loewenstein, George Ted O'Donoghue, and Matthew Rabin. 2003. “Projection Bias in Predicting Future Utility,” *Quarterly Journal of Economics* 118 (4): 1209–1248.

Malmendier, Ulrike and Stefan Nagel. 2011. “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?" *Quarterly Journal of Economics* 126(1): 373 -- 416.

Malmendier, Ulrike and Stefan Nagel. 2016. “Learning from Inflation Experiences." *Quarterly Journal of Economics* 131(1): 53-87.

Monroe, Kent. 1973. “Buyers' Subjective Perceptions of Price,” *Journal of Marketing Research* 10 (1): 70-80.

Mullainathan, Sendhil. 2000. “Thinking through Categories,” working paper.

Mullainathan, Sendhil. 2002. “Memory-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 117 (3): 735-774.

Mussweiler, Thomas and Fritz Strack 2001. “Considering The Impossible: Explaining The Effects of Implausible Anchors," *Social Cognition* 19 (2): 145 – 160.

Mussweiler, Thomas. 2003. “Comparison Processes in Social Judgment: Mechanisms and Consequences.” Psychological Review 110 (3): 472 – 489.

Nagel, Stefan and Zhengyang Xu. 2018. “Asset Pricing with Fading Memory,” working paper.

Nakamura Emi and Jon Steinsson. 2008. “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics* 123 (4): 1415–1464.

Pauwels, Kown, Shuba Srinivasan, Philip Franses. 2007. “When Do Price Thresholds Matter in Retail Categories?” Marketing Science 26 (1): 83 –100.

Pelham, Brett and Jeff Wachsmuth. 1995. "The Waxing and Waning of the Social Self: Assimilation and Contrast in Social Comparison,” Journal of Personality and Social Psychology 69(5): 825–838.

Piccione, Michele, and Ariel Rubinstein. 1997. "On the interpretation of decision problems with imperfect recall." *Games and Economic Behavior* 20(1): 3-24.

Rabin, Matthew. 1993. “Incorporating Fairness into Game Theory and Economics,” *American Economic Review* 83 (5): 1281-1302.

Rubinstein, Ariel. 1998. *Modeling bounded rationality*. MIT press, Cambridge MA.

Schwartzstein, Joshua. 2014. “Selective Attention and Learning.” *Journal of the European Economic Association* 12 (6): 1423–1452.

Sherif, Muzafer, Taub, Daniel and Hovland, Carl. 1958. “Assimilation and Contrast Effects of Anchoring Stimuli on Judgments” *Journal of Experimental Psychology,* 55(2): 150-155.

Simonsohn, Uri and George Loewenstein, 2006. “Mistake #37: The Effect of Previously Encountered Prices on Current Housing Demand." *Economic Journal* (116)508: 175-199.

Simonsohn, Uri. 2006. "New Yorkers commute more everywhere: contrast effects in the field." *Review of Economics and Statistics* 88(1): 1-9.

Simonson, Itamar, and Amos Tversky. 1992. “Choice in Context: Tradeoff Contrast and Extremeness Aversion." *Journal of Marketing Research* 29 (3): 281-295.

Sims, Christofer. 2003. “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3), 665–690.

Soman, Dilip and John Gourville. 2001. “Transaction Decoupling: How Price Bundling Affects the Decision to Consumer,” *Journal of Marketing Research* 38 (1): 30 – 44.

Stapel, Diederik and Jerry Suls. 2011. *Assimilation and Contrast in Social Psychology*. Psychology Press, New York, NY.

Taubinsky, Dmitry. 2014. “From Intentions to Actions: A Model and Experimental Evidence of Inattentive Choice." Mimeo, Harvard University.

Thaler, Richard. 1985. “Mental Accounting and Consumer Choice”, *Marketing Science*, 4 (3): 199 – 214.

Torgerson, Warren. 1958. *Theory and Methods of Scaling*. John Wiley & Sons, Inc, New York, NY.

Tversky, Amos and Daniel Kahneman. 1973. “Availability: A Heuristic for Judging Frequency and Probability." *Cognitive Psychology* 5(2): 207-232.

Tversky, Amos. 1977. “Features of Similarity.” *Psychological Review* 84(4): 327-352.

Niedrich, Ronald, Subhash Sharma and Douglas Wedell. 2001. “Reference Price and Price Perceptions: A Comparison of Alternative Models,” *Journal of Consumer Research* 28 (3): 339-354.

Wilson, Andrea. 2014. "Bounded memory and biases in information processing." *Econometrica* 82(6): 2257-2294.

Woodford, Michael. 2012. “Prospect Theory as Efficient Perception Distortion” *American Economic Review* 102 (3): 41 – 46.