

Saliency

Pedro Bordalo, Nicola Gennaioli, Andrei Shleifer¹

December 2021

Abstract

We review the fast-growing work on saliency and economic behavior. Psychological research shows that salient stimuli attract human attention “bottom up” due to their high contrast with surroundings, their surprising nature relative to recalled experiences, or their prominence. The Bordalo, Gennaioli and Shleifer (2012, 2013, 2020) models of saliency show how bottom up attention can distort economic choice by distracting decision makers from their immediate goals or from relevant choice attributes. This approach unifies probability weighting, menu effects, reference points, and framing as distinct manifestations of bottom up attention. We highlight new predictions and discuss open conceptual questions, as well as potential applications in finance, industrial organization, advertising, and politics.

Keywords: Attention, Memory, Reminders, Risky Choice

1. Introduction

¹ The authors are from Oxford University, Bocconi University, and Harvard University, respectively. They are grateful to John Conlon, Ben Enke, Sam Gershman, Thomas Graeber, Spencer Kwon, Giacomo Lanzani, and Josh Schwartzstein for helpful comments. Giovanni Burro provided outstanding research assistance. When citing this paper, please use the following: Bordalo, P, Gennaioli N, Shleifer A. 2021. Saliency. *Annu. Rev. Econ.* 14: Submitted. DOI: <https://doi.org/10.1146/annurev-economics-051520-011616>.

In psychology, a stimulus is salient when it attracts the decision maker's attention "bottom up," automatically and involuntarily. Automatic attention to salient stimuli has fundamental survival benefits, such as noticing and avoiding a barking dog. However, because the salience of a stimulus may differ from its current decision value, it can distract us from our goals and distort decisions. In standard economics, attention is either unlimited or, if costly, optimally deployed "top down" given current goals and expectations. This approach is useful, but does not recognize that goals often compete against bottom up stimulus-driven attention. Sometimes goals lose, even in important decisions.

Bottom up attention is crucial because markets abound with salient stimuli. The creative design of a Prada handbag contrasts with that of normal handbags, the high price of bottled water on a first visit to an airport is surprising, and the joy of driving a convertible on a sunny day is highly prominent. Salient stimuli attract the consumer's attention, possibly at the expense of his original goals. The design of the Prada bag draws attention away from its price, causing the consumer to spend more than originally planned on the bag. The surprisingly high price of water at the airport draws attention away from the benefit of drinking, causing the consumer to stay thirsty. The pleasure of driving a convertible on a sunny day leads the consumer to neglect driving in the snow, and to buy the car anyhow.

Salience can be strategically manipulated by firms or politicians. Salient ads or messages increase the prominence of some attributes while distracting consumers or voters from others. One example is the use of Twitter by former U.S. President Donald Trump. As Lewandowsky et al. (2020) show, in response to bad news in the Mueller investigation of Russia's interference in the Presidential election, Trump would immediately send provocative tweets on unrelated topics, such as China, jobs, or immigration. These "red herring" tweets would successfully capture the attention of even centrist and left-leaning news outlets, such as ABC News and the New York Times, distracting them from Muller's findings.

In recent years a growing literature has studied the role of salience in economic choice, both empirically and theoretically. We review this literature, including theory, experiments, and field evidence. We have three goals.

First, we describe the psychology of attention. Bottom up attention is automatically driven to stimuli that are contrasting, surprising, or prominent. Stimulus salience can then distract the decision maker, hindering his performance on various tasks. We show that this perspective brings common ground to economic analysis, which has used the term salience in different, often confusing, ways.

Second, we describe formal models of bottom up attention to contrasting, surprising, and prominent stimuli proposed by Bordalo, Gennaioli and Shleifer (2012, 2013, 2020), and show how they offer a way to unify many choice instabilities. Consider some examples.

Bottom up attention to lottery payoffs with high contrast generates overpricing of small stakes insurance and of right skewed stocks; attention to a good's attributes that are in stark contrast to those of other, even inferior, goods can account for the decoy effect.

Attention to surprising asset returns implies time varying risk attitudes in financial markets; attention to surprising prices can explain pricing strategies such as misleading sales.

Bottom up attention to prominent events can explain risk avoidance after hurricanes or earthquakes, but also risk taking as the memory of these events recedes. Bottom up attention to prominent product attributes can explain why consumers buy convertibles on sunny days, invest in financial products when their high returns (but not risk) are advertised, and neglect hidden attributes such as taxes or maintenance costs. Prominence can also account for framing effects. Describing purchases in terms of add-on pricing or in isolation, public health policies in terms of lives saved or lives lost, or the dictator game in terms of giving or taking, alter the salience of different features of the choice options, and hence their valuation.

In sum, many choice instabilities typically explained using a smorgasbord of ingredients such as probability weighting, reference points, menu effects, and exotic preferences can be viewed as the product of a well-documented and fundamental psychological force: bottom up attention to contrasting, surprising, and prominent stimuli.

Our third goal is to connect bottom up attention to research on reminders, nudges, and policy interventions. This work presumes that priming consumers with desirable goals or providing information improves choices by relaxing attentional or informational constraints. Bottom up attention highlights the limits of such designs. Priming, reminding, and informing consumers makes some aspects of their choice salient, and can have an excessive impact. A salesman truthfully informing a buyer that a smartphone can break, and selling him overpriced insurance, makes the buyer worse off. A Trump tweet, even if truthful, distracts voters from staying informed. An accurate statement that a stock looks like an early Google promotes bad investment. Much advertising provides true facts about some aspects of choice and distracts from others. As such, it can easily make the consumer worse off.

The paper is organized as follows. In section 2, we summarize the psychology of bottom up attention. In section 3, we present our approach to modeling it in economics. In section 4, we show how bottom up attention can account for a range of puzzling evidence in choice under risk. Section 5 presents some of the findings in consumer choice. Section 6 discusses research on reminders. Section 7 concludes and discusses avenues for future work.

2. The Psychology of Bottom Up Attention

William James (1890) famously described attention as “taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence”. This quote identifies two key aspects of attention, which remain center stage in current neuro-

scientific work (Nobre and Mesulam 2014). First, attention pertains to conscious experience of sensory perception *and* thought. We may attend to a stain on the wall or to memories of a vacation from last year. Second, we cannot consciously process all stimuli that reach our senses nor all information stored in memory. Attention is limited and therefore selective: the mind selects what we focus on, and downplays the rest.

Psychologists describe two broad selection mechanisms: top down and bottom up (also called voluntary and involuntary, or endogenous and exogenous, respectively). Top down mechanisms reflect higher level processes and in particular motivational factors such as the goals of the decision maker. We exhibit a remarkable ability to focus on the most task-relevant stimuli: we deliberately search for our car keys in a messy drawer, or keep track of our own kids in a crowded playground, shutting out other stimuli with lower perceived relevance. Models of rational inattention (Sims 2003; Woodford 2012, 2020; Khaw et al., 2021) and its cousin sparsity (Gabaix 2014) formalize this mechanism in economic decision-making. Attention may be allocated using incorrect beliefs, and as such distort decisions (Hanna et al., 2014; Schwartzstein 2014, Gagnon-Bartsch, Rabin, and Schwartzstein 2018), but it is still optimized given these beliefs *ex-ante*.

Yet current goals, while important, are not the only drivers. Attention is also automatically drawn to stimuli that are salient in a given context. While reaching for a drink at a cocktail party, a red jacket among black tuxedos spontaneously attracts attention “bottom up”, distracting us from our immediate goal. To quickly spot the red jacket, our brain must engage in some pre-attentive parallel processing of all jackets, and have an unconscious selection mechanism for the red one. In fact, our brain routinely engages in such pre-attentive parallel processing (Shiffrin and Gardner 1972) and is endowed with a system of “bottom up”

filters for salient stimuli that compete with goals in directing our attention (Treisman and Gelade 1980).² Saliency is the property of a stimulus that draws attention bottom up.

Recent work in neuroscience explores the neural architecture of top down and bottom up attention (Rao and Ballard 1999, Nobre and Mesulam 2014). Here we focus on two questions of more immediate interest to economics: What makes a stimulus salient? And how does saliency affect behavior?

On the first question, experimental evidence points to three factors that materially (though not exclusively) shape saliency: i) contrast with surroundings, ii) surprise, and iii) prominence. These factors have been studied for sensorial (visual and auditory) stimuli, but a stimulus may be salient also due to its emotional properties (Pashler 1998).

In a striking example of contrast, Treisman and Gelade (1980) show that subjects very quickly identify a red T target among blue L distractors, and that speed does not significantly decline as the number of distractors is raised. When a stimulus has high contrast with its surroundings, in this example in terms of color *and* shape, it quickly attracts attention, as if there is flawless parallel processing of all stimuli. When no stimulus is especially contrasting, such as when the same red T target occurs among blue T and red L distractors, attention is deployed serially, guided by goals, and identification takes time. Computational neuroscientists (Itti, Koch, and Niebur 1998) have built saliency maps for predicting the location of spontaneous human gaze in a visual scene, absent any goals. They successfully predict that gaze is directed to areas that display high contrast in one or more visual features such as color, shape, or orientation.

² Starting from the 1950s there has been debate between “Early Selection” theories (Broadbent 1958), according to which processing is serial, and “Late Selection” theories (Deutsch and Deutsch 1963) according to which there is full pre-attentive parallel processing of stimuli, so that selection occurs during conscious awareness. The evidence indicates that pre-attentive parallel processing, while not full (Kahneman and Chajczyk 1983), is substantial (Pashler 1998). Desimone and Duncan (1995) offer an influential theory of visual stimulus selection in which multiple stimuli are recorded by neurons and compete for identification. Attention then biases such competition in favor of the stimuli that are bottom up more prominent or more goal relevant.

Consider surprise next. We are highly sensitive to changes in sensory inputs, such as sudden changes in a visual scene (Liesefeld et al. 2017). Surprises relative to expectations or norms attract attention (Kahneman and Miller 1986): when opening a door, our attention is suddenly drawn to the doorknob if it is colder than usual (Hawkins and Blakeslee 2004, Clark 2013). The role of surprise underscores a point that we stress throughout: memory plays a key role in bottom up attention. In fact, surprise can be viewed as the contrast between a stimulus and the memories it evokes. Itti and Baldi (2009) enrich the visual saliency map with surprise relative to previous experiences (which they refer to as Bayesian surprise), and show that it also helps predict human gaze.

Finally, prominence refers to the idea that stimuli highly available to our senses or in memory are more salient. In visual perception, stimuli that are centrally located in the visual field are more easily detected (Peters and Itti 2007). Stimuli or locations that have recently attracted attention continue to do so, even if they are no longer task relevant (Remington, Johnston and Yantis 1992): after discussing whether families are having fewer children, we notice many babies on the street. Unlike contrast or surprise, which come from the comparison of stimuli, prominence comes from factors exogenous to the stimulus itself. An extensively advertised product is prominent in the consumer's mind, attracting attention bottom up. Visually prominent product attributes can play a key role in economic decisions.

Saliency affects behavior because salient stimuli are overweighed while non-salient stimuli are underweighted (Taylor and Thompson 1982). This is often desirable: it enables us to react swiftly when conditions change, and to monitor potentially relevant information for goals that are not yet explicit but may become so, such as sensing approaching danger. But when stimulus saliency and goal relevance are not aligned, as when surprise or suspense are used to direct our attention to an advertised product, bottom up attention interferes with our

goals and causes suboptimal decisions.³ This crucial implication is supported by a wide range of evidence. For example, performance in visual search tasks (measured by detection rates of targets or response times) deteriorates when a salient distractor, such as a red object against a field of green objects, is added to the image (Pashler 1988), or when a flicker occurs in a position away from the targets prior to display (Remington, Johnston and Yantis 1992).

Effects like these can also be created by surprise. We can focus on our reading despite background construction noise, but a sudden rise in noise levels distracts us from the book. For sensory prominence, Kahneman (2011) discusses many choice anomalies due to the “what you see is all there is” (WYSIATI) principle (Enke 2020). The features of a problem that are prominent or explicit receive a large weight in decisions, and cause underweighting or neglect of aspects not mentioned. These examples share a powerful thread: salient stimuli attract attention bottom up, interfering with a balanced weighting of our goals or interests.⁴

In sum, psychology highlights two important points. First, attention is often spontaneously allocated bottom up to contrasting, surprising, or prominent stimuli. Second, stimulus salience competes with goal relevance to capture our attention and sometimes wins, causing us to underweight or neglect initial goals.

3. Salience and Economic Choice

Bordalo, Gennaioli, and Shleifer (2012, 2013, 2020, henceforth BGS) present models of salience and bottom up attention in economic choice. The three papers consider different

³ Another way to say this is that bottom up attention is a general purpose tool that has evolved to achieve critical goals such as survival. However, it need not maximize utility in all decisions, including many important ones. The resource rationality approach (e.g. Lieder and Griffiths 2020) seeks to offer evolutionary foundations for “heuristics”, which are viewed as maximizing long run utility subject to the brain’s computational constraints. For instance, Lieder, Hsu and Griffiths (2014) argue that the high availability of extreme events can be justified by the brain’s recourse to importance sampling. It remains to be seen whether this or other approaches can offer a foundation for the joint role of contrast, surprise and prominence in attention and decisions.

⁴ The distracting role of bottom up factors can be reduced through learning. For instance, a person abruptly appearing on a sidewalk usually does not distract an experienced driver from his task. In other settings, the task is so engrossing that it interferes with stimuli that are potentially salient, as in the case of experimental subjects not noticing a gorilla when they count ball throws (Simons and Chabris 1999).

settings, but share a broad approach. When decision makers (DMs) choose, the attributes of choice options act as stimuli that trigger selective recall of normal attribute values from memory. The attributes of an option are then differentially salient based on: i) their contrast with the attributes of the other options, ii) the extent to which they are surprising compared to retrieved normal values, and iii) the prominence with which they are displayed or retrieved. The DM's attention is allocated bottom up to salient attributes, which are then overweighed, while non salient attributes are underweighted. Such underweighting reflects distraction.

This approach formalizes the three channels of bottom up attention and their impact on choice. It explains how valuation may depend not just on objective payoffs, but also on how the problem is presented or described. As we show in Sections 4, 5 and 6, this approach unifies many hitherto unrelated forms of choice instability.

We adapt the original models in BGS (2012, 2013 and 2020) to illustrate the three salience channels in the same setup. A good consists of $K > 1$ attributes (a_1, \dots, a_K) , where $a_k \geq 0$ is the hedonic value of attribute $k = 1, \dots, K$. The intrinsic valuation of the good is:

$$V = \sum_k \pi_k a_k, \quad (1)$$

where π_k is the decision weight attached to attribute k .

In riskless choice, such as buying a known bottle of wine, one can think of $K = 2$ attributes, quality a_1 and price a_2 , with price carrying a negative weight $\pi_2 = -1$ and quality (measured in dollars) a positive weight $\pi_1 > 0$. For a risky lottery, the attributes a_i are the absolute payoff values received in the possible states of the world and each payoff is weighted by its probability (with a negative sign if the payoff is negative). There are also hybrid cases like financial products, which have both stochastic (returns) and deterministic (fees) attributes. One could think of π_i as the optimal weights based on top down factors such as the DM's preferences or objective probabilities. If attribute k is salient, it draws attention and gets overweighed, distorting its decision weights into $\hat{\pi}_k > \pi_k$. If k is non-salient, the

consumer is distracted and underweights it, $\hat{\pi}_k < \pi_k$. Contrast, surprise, and prominence entail different mechanisms for distorting the weights, which apply in different situations.

To illustrate, consider a consumer choosing between two laptops, M and O. In the simplest case, the consumer sees all the relevant attributes (e.g., prices, processors, design of the laptops), and these attributes take the same values as they did in the past. In this case, attribute salience affects $\hat{\pi}_k$ and choice through *contrast* only. In a more complex situation, the consumer sees all attributes but the values of some, say price, differ from what the consumer has experienced in the past. In this case, attributes salience affects $\hat{\pi}_k$ and choice also through *surprise*. In a yet more complex case, the consumer sees some attributes such as design, processors, and prices, but does not see others such as taxes or maintenance costs. In this case, attribute salience can affect $\hat{\pi}_k$ and choice also via *prominence*. We describe these mechanisms one at a time, starting with the simplest of the three.

3.1 Contrast

Contrast captures the idea that a specific attribute of a good may “stand out” when the good is compared to alternatives. For instance, the more expensive laptop M may have a much more attractive design than O. Just as contrasting visual features attract gaze, the contrasting design attracts the consumer’s attention. Design is an objectively desirable attribute but, because it draws bottom up attention, it distracts the consumer from paying enough attention to other attributes such as price, biasing choice in favor of M.

To model contrast, BGS (2012) introduce a real valued, symmetric and bounded salience function $\sigma(a_k, \bar{a}_k) \geq 0$ measuring the contrast of the good’s attribute a_k relative to the average value \bar{a}_k of the same attribute for all goods in the choice set C . If the choice set C varies across situations, the salience of a_k for the good would also vary. Bottom up salience is context dependent. The salience function is characterized as follows.

Definition 1. *The salience of attribute k for a good in a choice set C is measured by a bounded function $\sigma(\cdot, \cdot) \geq 0$ that satisfies:*

- i) *Ordering: $\sigma(a_k + \mu\epsilon, \bar{a}_k - \mu\epsilon') > \sigma(a_k, \bar{a}_k)$ for $\mu = \text{sgn}(a_k - \bar{a}_k)$, $\epsilon, \epsilon' > 0$*
- ii) *Diminishing Sensitivity: $\sigma(a_k, \bar{a}_k) > \sigma(a_k + \epsilon, \bar{a}_k + \epsilon)$ for $\epsilon > 0$.*

According to ordering, attribute k is more salient for a good when its value a_k for that good is more different from the average value \bar{a}_k in the choice set. Just as a red dot is salient in a field of green dots but not in a field of identical red dots, the \$1000 price of laptop M is more salient if the price of laptop O is \$500 than if it is \$1000. Ordering generates an attentional externality: one good's attribute draws more attention if the same attribute in other goods becomes more different.

Diminishing sensitivity reflects the Weber-Fechner law of sensory perception: a given attribute difference is more salient at lower attribute values. Just as we perceive better a given increase in luminance against a dark background (stars are visible only at night), a \$100 price difference is more salient if the price level is \$400 than if it is \$1000. The diminishing sensitivity property of salience reflects a link between perception and attention: easier perception of payoff differences directs attention toward them. Itti and Koch (2001) incorporate diminishing sensitivity through stimulus normalization in their saliency map.

Ordering and diminishing sensitivity can pull in different directions, and specific salience functions pin down this balance. The homogeneous of degree zero function $\sigma(a, \bar{a}) = \frac{|a - \bar{a}|}{|a + \bar{a}|}$, which measures proportional differences between attributes, is a particular example that BGS often use for tractability. The key implications of salience hold more generally, and future work may find ways to measure the salience function.

Given a salience function, the intrinsic valuation (1) is distorted according to:

$$\hat{V} = \sum_k w_k \pi_k a_k, \quad (2)$$

where $w_k \equiv w(\sigma_k; \sigma_{-k}) \geq 0$ is a weighting function capturing bottom up attention to (good specific) salient attributes. The distorted attention weight is then $\hat{\pi}_k = w_k \pi_k$. Here σ_k denotes the salience of attribute k for the good, and the vector $\sigma_{-k} = \{\sigma_i\}_{i \neq k}$ denotes the salience of the good's other attributes. Weighting is formalized as follows.

Definition 2. *The weight w_k attached to the good's attribute k in \mathcal{C} increases in the salience σ_k of attribute k and weakly decreases in the salience of other attributes σ_{-k} .*

A choice set rendering a good's attribute more salient causes the DM to attach a higher weight to that good's attribute, in that choice set. It also causes the DM to attach a lower weight to attributes that, in the same choice set, are not salient. BGS offer different formalizations of weights, which however all capture this general idea. BGS (2012, 2013) assume weight normalization, $\sum_k \hat{\pi}_k = \sum_k \pi_k$. They further specify $w_k = \frac{\delta^{r_k-1}}{\sum_k \delta^{r_k-1} \pi_k}$, where $r_k \in \{1, \dots, K\}$ is the salience ranking of attribute k .⁵ BGS (2020) abstract from normalization and rank, and assume $w_k = \sigma_k$, which is continuous in salience. Normalization tends to amplify the effect of salience on valuation for it implies not only that a salient attribute attracts attention to itself, but also that it reduces attention to less salient attributes.

Context dependent bottom up attention to contrasting attributes offers a foundation for a form of choice instability typically explained with probability weighting, decoy effects and other choice set effects. We describe this in detail in sections 4 and 5.

The valuation formula of Equation (1) can be obtained from choice axioms. Ellis and Masatlioglu (2021) show that the rank-based specification of the salience model for riskless choice is a special case of a model of categorical valuation in which preferences follow standard choice axioms within categories of goods but display reference dependence across categories. Lanzani (2021) shows that the salience model for univariate lotteries obtains from

⁵ Lower ranking r_k means higher salience and larger weight w_k . Kontek (2016) shows that one pitfall of this rank dependent specification is that it creates discontinuities in valuation such that the certainty equivalent of a lottery is not always defined. BGS (2013) also offer a continuous formulation for decision weights.

a relaxation of transitivity, which captures the centrality of the ordering property. Köster (2021) generalizes the techniques of Lanzani (2021) to the choice between multivariate random variables. These papers build on the formalism of conjoint measurement of preferences (Fishburn 1989).

Other proposals also incorporate contrast-based distortions in choice. Koszegi and Szeidl (2013) present a model of focusing in which people overweight the attribute with the largest range in the choice set. Dertwinkel-Kalt et al. (2021) offer supporting evidence for that model's predictions in intertemporal choice. Bushong et al. (2020) capture the opposite idea that differences along the attribute with the largest range are underweighted due to perceptual normalization and offer supporting evidence for this effect. The model we just presented captures both contrast and diminishing sensitivity to attribute values, and allows for different attributes to be salient for different goods. Crucially, other factors also shape bottom up attention, and expand the reach of the model.

3.2 Surprise

Temporal variation in the values of some attributes, such as price, opens room for surprise. When the consumer sees the price of laptop M, he retrieves from memory past prices he has seen for the same or similar laptops. A current price of M that is much higher than its normal value draws attention bottom up. Again, price is relevant for the DM's goals, but excessive attention to it causes an over-reaction, distracting the consumer from other relevant attributes, and biasing choice away from M. This occurs even if recalled experiences convey no information, as in the "background contrast" effects reviewed in Sections 4 and 5.

To capture surprise, BGS (2020) introduce memory into the model. A good's attributes (a_1, \dots, a_K) act as cues that trigger retrieval of the attribute values of the same good that the DM has experienced in the past. Such retrieval gives rise to a "normal" (average)

version of the good, denoted by (a_1^n, \dots, a_K^n) , which is individual specific. Averaging the retrieved norms for all goods in the choice set C , the DM forms a memory-based reference $(\bar{a}_1^n, \dots, \bar{a}_K^n)$. The salience of attribute k is then given by $\sigma(a_k, \bar{a}_k^n)$: the DM overweighs the attribute that is most surprising in having the highest contrast from its normal value.

To formalize the consumer's recall process, BGS (2020) build on three well-established regularities in human recall: frequency, similarity and interference. When cued with a piece of data, we tend to recall experiences that are frequent and similar to it, which interferes with recall of less frequent and less similar ones. For instance, when thinking about "white things from the kitchen," we are likely to recall milk, given that it is experienced every day. In turn, this makes us more likely to recall "yogurt" or "cheese," and less likely to recall "plates," because the more similar dairy products interfere. In this vein, a choice option reminds a DM of past experiences with similar goods, which are used to compute normal attribute values. As we discuss in Section 5, regularities of human memory help account for the different reference points used in behavioral economics, and make new predictions for when the reference point is not the rational prior.

BGS (2013, 2020) show how memory based surprise distorts choice. Suppose that when choosing between laptops M and O the consumer retrieves normal prices of \$1000 and \$800 for the laptops, leading to a memory-based reference price $\bar{p}^n = \$900$. If the current prices are $p_M = \$1200$ and $p_O = \$800$, then M's price is surprisingly high and draws attention. Thus, the consumer overweighs p_M and displays a relative neglect for the quality of M, reducing his valuation for this laptop. The valuation of O is instead unaffected because the price of this laptop is normal. When seeing the high current price, even a rational consumer would reduce his valuation of M compared to the normal case $p_M = \$1000$. However,

salience *also* boosts the consumer's weight on the price of M, rendering him more price elastic. This mechanism can account for several puzzles in consumer behavior.⁶

Surprise relative to a memory based norm offers a foundation for a second source of choice instability: reference point effects, as we show in Sections 4 and 5.

3.3 Prominence

Product attributes relevant to choice are sometimes not observed or not displayed prominently, as with fine print. In a rational world, the consumer would think of all sufficiently relevant attributes, and possibly assume the worst for those he is uncertain about that are not mentioned or advertised. In reality, the psychology of selective memory suggests that visible attributes draw attention to themselves, which may interfere with retrieval of goal relevant yet hidden attributes, as in Kahneman's (2011) WYSIATI principle.

BGS (2020) use their model of memory to formalize this notion.⁷ Suppose that the consumer is cued with information $\kappa = (a_k)_{k \in P}$ that captures the good's attributes that are naturally available to perception or advertised. The remaining attributes $(a_k)_{k \in I}$ are not shown. The consumer retrieves from his database norms $(a_k^n)_{k \in P}$ for the visible attributes, but may also spontaneously recall experiences with hidden attributes. Suppose that invisible attribute $k \in I$ is recalled with probability r_k . Valuation is then on average given by:

$$\hat{V}_P = \sum_{k \in P} w_k \pi_k a_k + \sum_{k \in I} r_k \tilde{w}_k \pi_k a_k^n. \quad (3)$$

The first term captures the valuation of visible attributes $k \in P$, whose weight w_k is distorted by contrast and by surprise relative to their retrieved norm a_k^n as in Section 3.2. The

⁶ In this model, surprisingly low prices render prices salient, tilting choice toward cheaper goods. At the extreme, when the good is not very valuable, this may cause the consumer not to buy. BGS (2020) propose a formalization of surprise avoiding this prediction. When considering a single good, the valuation of each attribute is anchored to the norm and adjusted based on the surprise $(a_k - a_k^n)$ times its salience $\sigma_k(a_k, a_k^n)$:

$$V = \sum_k \pi_k [a_k^n + \sigma_k(a_k, a_k^n) * (a_k - a_k^n)],$$

so that surprising low prices increase valuation, favoring buying over not buying.

⁷ A first draft of that paper (2015) applied the formalism presented below to shrouded attributes and reminders.

second term captures the valuation of invisible attributes: these tend to be underweighted on average because they may fail to be recalled, $r_k < 1$, and their weights \tilde{w}_k only depend on contrast, not surprise (because the actual values of these attributes are not observed).

Equation (3) implies that prominent features attract attention bottom up, especially if they are contrasting or surprising, while non-prominent features – albeit goal relevant – are non-salient, and can thus be neglected. This is relevant to advertising: making some features of laptop M prominent, such as its design, may interfere with recalling its high maintenance costs, causing the DM to underweight them and biasing choice in favor of M.

Prominence is a foundation for a third source of choice instability, framing effects, whereby decisions depend on normatively irrelevant changes in the description of a problem.

In sum, bottom up attention to salient stimuli offers a natural foundation for three major sources of preference instability: choice set effects (contrast), reference point effects (surprise), and framing (prominence). We next discuss some relevant evidence, as well as new predictions that follow from the psychology of bottom up attention and memory.

4. Choice under Risk

Over the past decades, social scientists have documented significant instability in risk attitudes in both the lab and the field. This instability is inconsistent with Expected Utility Theory (EUT) and has motivated alternative frameworks, most notably Prospect Theory (PT, Kahneman and Tversky 1979, Tversky and Kahneman 1992). Bottom up shifts of attention across lottery payoffs offer an explanation: Subjects are risk seeking when a lottery's upside is salient, and risk averse when its downside is salient. Changes in the salience of lottery payoffs caused by the payoffs of other available lotteries (contrast), past experiences (surprise), and contextual cues (prominence), explain different forms of instability documented in the lab and the field and entail novel predictions for behavior towards risk.

4.1 Contrast and Risky Choice

A key manifestation of unstable risk preferences is the tendency to take right skewed risks while avoiding left skewed ones (Kahneman and Tversky 1979). Consumers pay too much to reduce deductibles in home and auto insurance (Sydnor 2010, Barseghyan et al. 2013), or to buy extended warranties (Abito and Salant 2019). Overweighting of unlikely but large losses, and not just loss aversion, is necessary to account for this evidence. Consumers also pay too much to buy state lotteries and right skewed stocks, suggesting overweighting of unlikely large gains (Barberis and Huang, 2008; Chiappori et al, 2019, Lockwood et al 2021).

Such instability naturally arises from bottom up attention driven by contrast. When a DM is offered a choice between a lottery $(L, p; 0, 1 - p)$ and its expected value $(pL, 1)$, the contrast between payoffs L and pL increases as the probability p falls. Payoff L becomes more salient as the lottery gets more skewed, while the zero payoff becomes less salient. As a result, payoff L draws the DM's attention and gets overweighted while the 0 payoff gets underweighted. This mechanism causes the same investor to be risk seeking with an unlikely large upside $L > 0$ and risk averse with an unlikely large downside $L < 0$.

In salience theory payoff contrast and hence choice can be shaped by changes that are wholly irrelevant for EUT, such as varying lotteries' common consequences, and even for Cumulative Prospect Theory, such as varying their correlation structure, or the choice set.

The Allais paradoxes represent a glaring example of such instability (Allais 1953). The choices in Panel A in Figure 1 differ from those in Panel B only by replacing, in both lotteries, a \$0 payoff with a \$2400 payoff in a 0.66 probability event. According to the "independence axiom" of EUT, such common payoff change should not affect choice. Yet, in Panel A subjects are risk averse, preferring the sure thing O_s to O_r , while in Panel B they are risk seeking, preferring the riskier lottery P_r to P_s (Kahneman and Tversky 1979).

Panel A	Panel B
$O_r = \begin{cases} \$2500 & \text{prob. } 0.33 \\ \$2400 & \text{prob. } 0.66 \\ \$0 & \text{prob. } .01 \end{cases}$	$P_r = \begin{cases} \$2500 & \text{prob. } 0.33 \\ \$0 & \text{prob. } 0.67 \end{cases}$
$O_s = \{ \$2400 \text{ prob. } 1$	$P_s = \begin{cases} \$2400 & \text{prob. } 0.34 \\ \$0 & \text{prob. } 0.66 \end{cases}$

Figure 1

Contrast-driven bottom up attention explains this behavior. In Panel A, the state in which lottery O_r pays 0 and lottery O_s pays 2400 is highly salient, which follows directly from ordering in Definition 1. Here bottom up attention focuses the DM on the risky lottery's downside of 0 and distracts him from the upside of 2500, causing risk aversion.

In Panel B, both lotteries face the same downside of zero, and the highest contrast state is the one in which lottery P_r pays 2500 and lottery P_s pays zero. Bottom up attention now focuses the DM on the risky lottery's upside of 2500 and distracts him from the slightly more likely downside of 0, inducing risk seeking behavior. This mechanism also accounts for the common ratio Allais paradoxes (BGS 2012).

To deal with the Allais paradoxes, Prospect Theory proposes that subjects mechanically overweight low probabilities and underweight high ones. Subjects in Panel A are thus risk averse because they are very sensitive to the 1% probability that O_r yields 0, but in Panel B they are insensitive to the same extra 1% probability of 0 with P_r compared to P_s .

Saliency theory endogenously produces a weighting function, but one in which distortions do not depend only on probability values. Due to bottom up attention shifts, payoffs matter: small probabilities are overweighted if and only if they are associated with extreme and hence salient payoffs.⁸ But the choice context matters too, in particular the

⁸ In this way, saliency naturally captures the rank-dependent properties of successive refinements of Prospect Theory (Tversky and Kahneman 1992). Rubinstein (1988) offers a related account in which subjects underweight the similar probabilities of lotteries P_r and P_s , thus overweighting their payoff differences.

correlation between lotteries. To see this, consider again the lotteries in Figure 1, Panel B, but suppose that they are correlated. There are then only three states of the world as below.

$$P_r = \begin{cases} \$2500 & \text{prob. } 0.33 \\ \$0 & \text{prob. } 0.66 \\ \$0 & \text{prob. } .01 \end{cases} \quad P_s = \begin{cases} \$2400 & \text{prob. } 0.33 \\ \$0 & \text{prob. } 0.66 \\ \$2400 & \text{prob. } .01 \end{cases}$$

Relative to Panel B, this correlation structure drastically affects salience: now it is the state in which P_r obtains 0 and P_s obtains 2400 that is most salient, as in Panel A. The DM focuses on this state and neglects the risky lottery's upside of 2500, yielding risk averse choice. Expressing *the very same* lottery of Panel B in this way leads preferences to change by changing bottom up attention.⁹ The Allais paradox should then disappear, because the same state is salient regardless of the common consequence. This is exactly what the evidence shows; see BGS (2012), Bruhin et al. (2020), and Frydman and Mormann (2020).

The probability weighting function of PT cannot account for this evidence: the correlation structure does not change the probability of each payoff. Using theory and experiments, Dertwinkel-Kalt and Koster (2020) show that correlation also affects skewness preferences: if two right skewed lotteries are perfectly correlated – so that they receive their upsides together – the salient state might be the downside of both. This causes a preference for the lottery with lower skewness. Bottom up attention shifts explain this instability.

Further evidence of the role of attention in risky choice is provided by Pachur et al (2018). In one experiment, subjects chose between several pairs of binary lotteries while their attention to outcomes and probabilities is measured.¹⁰ More attention to losses is associated with more loss-averse behavior. In another experiment, the experimenters exogenously varied the length of time that lottery gains or losses were shown to subjects.

⁹ The “common consequence” state where both lotteries yield \$0 is the least salient one, and cancels out in the DM's valuation. We thus recover the sure thing principle, while probabilistic sophistication fails.

¹⁰ Measures of real time attention allocation such as eye tracking or mouse tracking can also be used in lab or field experiments (Krajbich, Armel, and Rangel 2010, Mormann and Russo 2021). For a review of methods of process tracing, see Schulte-Mecklenbeck et al (2017).

Here again, attention to losses is associated with loss-averse behavior, but this time the evidence supports a causal link.

Payoff contrast shaping risk attitudes has profound implications for many settings. BGS (2013b) show theoretically that salience may shed light on several asset pricing puzzles: overpricing of right-skewed assets, such as growth stocks, a large equity premium, due to aversion to the aggregate market's left skewness, and countercyclical risk aversion.¹¹

4.2 Surprise and Choice Under Risk

An unusually large state lottery jackpot may entice someone to buy a ticket, despite the very low odds, because the prize is surprisingly high. An investment yielding 2% seems much less attractive after a period of 10% returns than when a 1% return is normal. Bottom up attention to surprises can account for these intuitions, unifying a host of evidence on choice under risk from both the lab and the field.

In an experiment by Soltani et al. (2012), subjects are shown three lotteries: a risky target L_r , a safe target L_s , and a decoy L_d . In a subsequent choice, L_d is removed and subjects choose between L_r and L_s . Behavior is consistent with a “phantom decoy” effect: if L_d dominates L_r by having the same win probability but a higher upside, subjects are more likely to choose the safe lottery L_s . Soltani et al. (2012) explain the findings with range normalization, similarly to Bushong et al. (2020). Salience offers an alternative account: in the choice stage, subjects retrieve a norm for the lottery upside that includes the upside of L_d . This renders the upside of L_r less surprising and hence less salient than in the absence of L_d . The DM focuses on the lower win probability of L_r and neglects its larger upside, becoming risk averse. Because these effects are due to memory and not choice alternatives, they cannot be explained by motivational theories such as regret-avoidance (Loomes and Sudgen 1982).

¹¹ Countercyclical risk aversion arises because at low payoff levels, lottery downsides are relatively more salient so people are more risk averse. BGS (2012) offers experimental evidence consistent with this prediction.

Lian et al. (2019) use experimental and field data to study investors' reach for yield. Investors are much more likely to choose a riskier asset delivering, say, a 6% expected return over a safe one yielding 1% after experiencing high rates of return for all assets in the recent past. This finding is consistent with the idea that retrieval of past high rates causes current rates to appear surprisingly low for all assets. This shifts the investor's attention to returns and causes him to neglect risk, favoring the high return-high risk option.

Memory-based surprise can also account for preference reversals from willingness to pay to choice. Lichtenstein and Slovic (1971), Grether and Plott (1979) and Tversky, Slovic, and Kahneman (1990) show in the gain domain that subjects may prefer a lottery L_s to a riskier one L_r and yet report a lower minimum selling price for L_s than for L_r . Such reversals can even violate "procedural invariance" (Tversky et al. 1990): subjects may report a price for L_r higher than its expected value, and yet prefer the expected value to L_r in binary choice.

BGS (2012) show that salience theory delivers preference reversals. When choosing between L_r and its expected value, subjects may focus on the lottery loss. Matters change when valuing L_r in isolation: the DM retrieves a norm of not having the lottery at all, which renders its upside surprising and its zero downside normal. This causes risk seeking.¹²

4.3 Prominence and Choice Under Risk

Prominence driven bottom up attention can account for instability in risk attitudes when: i) the description of payoffs changes, and ii) rare events are made more prominent.

In Kahneman and Tversky's (1979) Asian disease problem, DMs choose between two emergency plans to save 600 people at risk. In plan A, 200 people are saved for sure. In plan B there is a one third chance 600 people are saved and a two thirds chance 0 people are saved. The same choice can be described differently: under plan A, 400 people die for sure,

¹² This effect is strong provided the DMs' norm is a zero or low payoff. This mechanism can then be tested by priming DMs to recall norms at different payoff levels, e.g. by manipulating payoffs of lotteries seen earlier.

whereas under plan B there is a one third chance that 0 people die and a two thirds chance that 600 people die. Subjects tend to choose plan A when the choice is described in terms of lives saved, and plan B when it is described in terms of lives lost.

KT account for this evidence via a shift in the reference point, but they do not explain why such a shift would occur when only the description of payoffs changes. Bottom up attention helps explain this puzzle using the idea in Section 3.3 that prominent attributes interfere with the retrieval of hidden ones: when thinking about lives saved, subjects fail to consider the implications for lives lost and vice-versa, with drastic implications for valuation (Equation 3). In the “lives saved” version, the salient payoff of the risky option is “0 lives saved”, due to diminishing sensitivity. Attention to it causes risk aversion. By the same token, in the “lives lost” version the salient payoff of the risky option is “0 lives lost”. Attention to it prompts the subjects to gamble for resurrection.

The role of prominence is amply documented in the field, where payoff states are often hidden from decision makers, or fail to come to mind. Dessaint and Matray (2016) study how corporate managers respond to hurricane risk. When a hurricane strikes nearby, managers sharply raise cash holdings and express more concerns about hurricane risk, even though the actual risk remains unchanged. Bottom up attention offers an intuitive explanation: a prominent event cues salient risks, even when it does not provide new information. This mechanism causes an over-reaction inconsistent with rational learning, explaining why in the data cash holdings revert back to normal as the prominence of the risk subsides. Kunreuther et al (1978) similarly show that households rush to buy flood insurance after a flood, but again the effect is temporary. Shifting prominence of risks remedies a key difficulty with Prospect Theory’s mechanical overweighting of unlikely events: there are many possible disasters, but mostly we are blissfully unaware of them.

Shifts in prominence can also play a role in asset markets, where they could be driven by recent returns (as in models of time varying disaster risk) or by advertising. Célérier and Vallée (2017) study the European market for retail structured financial products. These products typically offer a high return under their best-case scenario – the headline rate – that is nested in a complex payoff formula, but is prominently advertised to investors. The authors find that when interest rates fall, the advertised headline rates rise, complexity increases, and risky products proliferate. These products are more profitable to banks distributing them. Saliency accounts for these findings: the advertised headline rate attracts attention and looks surprising to investors compared to the low safe rate. It also interferes with attention to downside risk or fees, increasing investor demand.

5. Consumer Choice

Substantial evidence from both the lab and the field shows systematic instability in consumer preferences resulting from goal-irrelevant changes in the choice environment. Bottom up shifts of attention to salient product attributes (BGS 2013, 2020) account for this behavior and yield new predictions, including for pricing strategies and market equilibria.

5.1 Contrast and Consumer Choice

According to the weak axiom of revealed preferences, adding a dominated option to a choice set should not affect which good is chosen. The “decoy effect” (Huber, Payne, and Puto 1982; Tversky and Simonson 1993) is a striking rejection of this axiom: when choosing between a good toaster for \$20 and a somewhat better one for \$30, most experimental subjects choose the cheaper toaster. But when a marginally superior toaster is added to the choice set for \$50, subjects switch to the middle toaster. There is evidence that such decoys are exploited by marketing strategies (Heath and Chatterjee 1995; Wu and Cosguner 2020).

Traditional explanations are based on loss aversion relative to a reference good (Tversky and Kahneman 1991; Tversky and Simonson 1993; Bodner and Prelec 1994).

In salience theory, choice instability created by decoys is a manifestation of contrast-driven bottom up attention. When choosing between two toasters, the \$30 price stands in high contrast to \$20. Attention is driven bottom up to price and away from quality, causing the consumer to refuse to pay \$10 extra for the better toaster. When the expensive \$50 toaster is added, the \$30 price is not as contrasting anymore. When assessing the middle toaster, attention now switches from price to quality, and the consumer is willing to pay the extra \$10 for it. Bottom up attention yields a novel prediction: the decoy cannot benefit the lower quality good. Indeed, adding a very low quality toaster makes the quality of the good toaster even more salient. Experimental tests confirm this prediction (Heath and Chatterjee 1995).¹³

This mechanism can also account for instabilities due to framing. In Savage (1954), a consumer prefers to pay \$17,500 for a car equipped with a radio to paying \$17,000 for a car without a radio, but would not buy a radio separately for \$500 after agreeing to buy a car for \$17,000.¹⁴ In salience theory, the radio's price is much more contrasting and hence salient when assessed in isolation than when it is added on top of the large cost of the car, due to diminishing sensitivity (BGS 2013). This effect can help explain why add-on pricing boosts demand (Soman and Gourville 2001), why the "latitude of acceptance" of a certain price grows with the price level (Koschate-Fischer and Wullner 2017), and also why price dispersion rises with a good's average price (Pratt, Wise, and Zeckhauser 1979).

In this example, a higher average market price reduces price contrast, which in turn reduces price elasticity of demand. BGS (2016) show how, via this effect, innovations reconfigure market competition. A major cost-cutting innovation leads to an equilibrium with

¹³ Bushong et al. (2020) offer a range normalization account of the decoy effect and compare it to BGS (2013). BGS (2013) show that salience can also help explain the compromise effect (Simonson 1989).

¹⁴ Also, subjects thinking of buying a calculator for \$15 and a jacket for \$125 are more likely to agree to travel for 10 minutes to save \$5 on the calculator than on the jacket (Kahneman and Tversky 1984; Kahneman 2011).

salient price, price sensitive consumers, and underprovided quality. A quality improving innovation leads to an equilibrium with salient quality, price insensitive consumers, and overprovided. Innovations affect choices of all consumers, not just the marginal ones, by switching salience from price to quality. BGS illustrate these predictions with commoditization of the airline market and decommoditization of the coffee market.

Framing and decoy effects can thus be unified: normatively irrelevant changes in the choice set or in its description affect attribute contrast, in turn affecting whether quality or price is salient, and thus changing valuation. The same mechanism accounts for the “background contrast effect”, which we discuss next.

5.2 Surprise and Consumer Choice

Surprise relative to past experiences creates salience effects in markets. A first-time traveler at an airport may refuse to buy bottled water for \$4, even if very thirsty, due to the sticker shock caused by the comparison with the \$1 price he usually pays at the supermarket. “Background contrast” experiments (Simonson and Tversky 1992) emulate the role of past experience by having subjects choose among similar goods in two stages. Subjects are more likely to choose a cheap good in the second stage if they saw lower prices first.

Surprise driven bottom-up attention accounts for such choice instability. Prices experienced in the first stage are recalled in the very similar second stage. If the current price is unusually high, the consumer is surprised. His attention focuses on price, away from quality, so he refuses to buy. This is similar to over-reacting to the sticker shock when first seeing the price of water at the airport. BGS (2013) show that the effect of price reductions may be asymmetric if the consumer is making a vertically differentiated choice between a high and a low quality good. In this case, an increase in the price of the high quality good

causes a sticker shock, but a price reduction may render the price closer to the price norm for the category. This shifts attention to the good's quality, raising demand.¹⁵

Similar effects are documented in the field. Simonsohn and Loewenstein (2006) find that movers from cities with expensive real estate such as San Francisco rent more expensive and larger apartments in cheaper destination cities such as Pittsburgh than do movers from low-rent locations. BGS (2019) revisit this evidence and show that the interplay of memory and bottom up attention yields new implications that find support in the data. Hastings and Shapiro (2013) show that parallel drops in the prices of different gas grades cause consumers to switch to premium gas, to an extent that cannot be accounted for by income effects.¹⁶

Bottom up attention shaped by surprised relative to a memory-based norm offers a foundation for reference point effects. Crucially, because norms are based on associative recall, they may be influenced by goal-irrelevant contextual cues, and change behavior by creating artificial surprises (BGS 2020). This yields a theory in which rational expectations (Koszegi and Rabin 2006), status quo, and adaptive reference points arise as special cases, and can be predicted on the basis of consumers' experiences and contextual cues. It also predicts that reference effects can be created by priming: a consumer considering whether to buy a bottle of wine costing \$60 at a restaurant may be persuaded not to buy by a reminder that the same bottle costs only \$30 at the store, even though he may well have bought the bottle in the absence of the reminder. Due to associative memory, reference points can depart from the rational prior, so that surprise may be artificially triggered by irrelevant contextual cues. This mechanism accounts for Thaler's (1985) famous beer experiment: a consumer

¹⁵ BGS (2020) account for the background contrast experiment for valuing individual goods. BGS (2013) consider the case of a vertically differentiated choice set. With vertical differentiation, a reduction in the price of the high quality good exerts non-monotonic effects on salience: a moderate price reduction reduces price salience, a large reduction increases it. A full reconciliation of BGS (2013) and (2020) is left for future work.

¹⁶ Many other papers find evidence of background contrast effects. Simonsohn (2006) shows that movers to a new city choose longer commutes the longer their commute in the origin city. Bhargava and Fisman (2014) show, in a speed dating website, that prior partner attractiveness reduces the likelihood of an affirmative mating decision. Hartzmark and Shue (2018) show, in financial markets, that large and positive earnings surprises reduces the announcement returns of other firms in subsequent days. Bottom up attention driven by surprise can account for these findings (see BGS 2020).

considering drinking a beer on the beach has a higher price norm, and thus a higher willingness to pay, if the beer is described as being purchased at a nearby resort (which triggers recall of high prices) than at a nearby shack (which triggers recall of low prices).

Firms can strategically manipulate bottom up attention by creating an artificial surprise. Ortmeyer, Quelch and Salmon (1991) show that the majority of revenues of large department stores comes from fairly frequent deep discounts, which is hard to reconcile with standard price discrimination. One explanation is that retailers inflate the regular price to lure consumers into buying during sales. Consistent with this explanation, several US states as well as the EU restrict fake sales. Saliency theory accounts for this practice: the high regular price is retrieved during the sale and acts as an expensive decoy, boosting consumers' willingness to pay. BGS (2013) also show the limits of this effect: it can only occur infrequently enough and with limited competition (which reduces price comparability).

The interaction of price contrast and price surprise in driving bottom up attention is a source of new predictions. Prices that are high relative to norms, such as a \$4 water bottle, attract attention, due to the ordering property of saliency, and render the consumer price sensitive. But prices that are normally high, such as wine in a restaurant, render given price differences less salient, due to diminishing sensitivity. The same consumer may then be price elastic for wine at a store, price inelastic for wine at a high-end restaurant, but become price elastic if at the high-end restaurant wine is unusually expensive. Dertwinkel-Kalt et al. (2017) find experimental support for this prediction, while carefully controlling for intrinsic quality.

5.3 Prominence and Consumer Choice

In many markets, only some of the attributes of goods are visible at the time of choice. Other important attributes may not be top of mind because they materialize in the

future, are not advertised, or are strategically obfuscated. Prominence-driven bottom up attention then creates choice instability that accounts for a range of evidence from the field.

The famous projection bias (Loewenstein, O'Donoghue, and Rabin 2003) illustrates these ideas. Conlin, O'Donoghue, and Vogelsang (2007) show that catalog orders of cold-weather items spike on very cold days, and items ordered on such days are more likely to be returned. Busse et al. (2015) find that consumers are more likely to buy convertible cars on sunny days, even if they have already owned a convertible in the past. Chang, Huang, and Wang (2018) find that on days when air pollution is high, there is a spike in the purchase of health insurance. Many contracts are canceled when air quality subsequently improves.

In these examples, current context causes one attribute of a good to be prominent: the warmth of a sweater, the joy of driving a convertible on a sunny day, or the risk of respiratory disease on a high pollution day. Such context acts as a cue for a good's value, and interferes with the recall of other attributes: limited storage space for sweaters at home, the pain of driving a convertible in the snow, or the cost of health insurance. As the prominence of the initial cues subsides, the consumer recalls the shortcomings of the goods he purchased. Initial choices are undone because attention to the upside was bottom up, not goal driven.

By the same mechanism, bottom up attention sheds light on Chetty et al.'s (2009) idea of "tax salience." They show that U.S. consumers tend to neglect sales taxes, but report them accurately when prompted to do so, which is symptomatic of selective recall. Taxes are characterized as "non-salient" because they are paid at the checkout. This effect reflects prominence: the salient cue at the moment of choice, when the consumer is in the aisle, is the pre-tax price. Attention is directed to the listed price, and away from taxes. Finkelstein (2009) relatedly shows that drivers underweight changes in highway tolls paid not at the tollbooth but electronically. The use of the same term "salience" for both shrouded taxes or fees and

for large lottery payoffs may seem confusing, but it is in fact fully consistent with bottom up attention, which does not just depend on contrast and surprise, but also on prominence.

A large literature documents significant distortions due to consumers' inattention to hidden attributes, such as the need for replacement cartridges when buying a printer (Gabaix and Laibson 2005). Bottom up attention may help understand why firms can successfully obfuscate product features. This can occur when the cue, here the printer, does not remind the consumer of the expensive cartridge, as is likely to happen for a first time buyer, but less so for a repeat buyer who had previously paid a memorable price for a replacement cartridge.

6. Salience and Welfare

Consumer neglect of options' attributes has spurred policy efforts aimed at manipulating the choice architecture through nudges and reminders (Thaler and Sunstein 2008, Mullainathan et al. 2012). The idea is that reminders help consumers optimally incorporate previously neglected information into their decisions (Bernheim and Rangel 2009). We call consumers who abide by this description Forgetful But Otherwise Rational, or FBOR. However, in the world of bottom up attention described in Section 3, reminders do not merely provide information. They may also: i) cause surprise, and ii) interfere with the recall of other attributes. We call consumers vulnerable to such bottom up attention shifts Forgetful And Salient Thinkers, or FAST. The distinction between FBORs and FASTs is useful for thinking about: i) the evidence on which reminders work and which ones do not, ii) the welfare effect of reminders and iii) strategic use of reminders by firms.

Consider first which reminders work. When given a generic reminder to save, people save a bit more, but reminding them of specific future goals such as "saving to pay for sickness" leads to a larger rise in savings rates (Karlan et al 2014). The salience of reminders can account for this. The generic reminder may cue a FAST to recall frequently experienced

benefits of savings, such as a larger bank account, but these recalled benefits may be less salient than the costs, such as foregoing a vacation. In this case, and perhaps pervasively with generic reminders, a FAST would underreact. Reminding a FAST to save for future emergencies is more effective because these are infrequent or never experienced, so they do not easily come to mind (Sussman and Alter 2012), and because – once top of mind – they contrast strongly with other experiences. A reminder about extreme but unlikely outcomes would have a small effect on FBORs but a large effect on FASTs. Akerlof (1991) and Handel and Schwartzstein (2018) offer a related take on these issues.

The distinction between FBOR and FASTs is even starker for welfare. Following Bernheim and Rangel (2009), a large body of work measures the welfare effect of reminders by the extent to which they change consumer behavior (Chetty et al 2009, Goldin and Homonoff 2013, Calzolari and Nardotto 2017, Karlan et al. 2014, Patrick et al. 2009, Stango and Zinman 2014). This approach is correct for FBORs. When reminded of a \$4 sales tax, an FBOR increases the perceived price of the good by \$4 and makes a better decision. Matters are different if the consumer is FAST. When informed about the \$4 sales tax, a FAST may find the good surprisingly expensive, causing him to overweight price and to cut demand more than an FBOR would. Cueing taxes may also interfere with recall of shrouded qualities of the good, causing an even larger overreaction. The change in the consumer's behavior may suggest a large improvement even though welfare may have decreased due to overreaction, similarly to buying a convertible on a sunny day.

Salience offers a way to assess when reminders may hurt consumers. This can be done in a setting where the reminder is displayed for more than one period and consumer behavior is observed throughout. Overreaction to surprise would then imply that the elasticity of

consumer behavior to the reminder is high in the short and lower in the long run, when the reminder becomes normal.¹⁷

The ambiguous welfare effect of reminders is evident in the strategic disclosure of information by firms, a form of persuasion that benefits firms rather than consumers.¹⁸ Suppose that a consumer is buying a new phone and the salesman offers him insurance against the phone breaking. In a world of FBORs, this information improves welfare. In a FAST world, the information may remind the consumer of the salient risk of phone damage, redirecting his attention toward that risk, and causing him to overpay for insurance, in line with the evidence in Section 4.¹⁹ Common marketing strategies can be viewed as ways to exploit the FASTs' overweighting of salient information. They change consumers' behavior by distracting them from other goals. With endogenous allocation of attention to salient stimuli, the assumption that information makes consumers better off, or that one can infer welfare gains solely from changes in behavior, is difficult to swallow.

7. Concluding Remarks and Open Issues.

Choice behavior is driven by what decision makers attend to. The standard economic approach assumes utility maximization, which is another way of saying that goal-relevant stimuli are top of mind. But attention is also drawn, in a bottom up fashion, to stimuli that are salient due to their contrast, surprise, and prominence. Because saliences changes across goal-irrelevant conditions, bottom up allocation of attention creates state-dependent distraction from relevant goals and thus choice instability. Contrast-driven bottom up

¹⁷ In contrast, noisy attention to individual reminders would imply a stronger long run vs short run adjustment.

¹⁸ Studies of such persuasion that are not linked to salience or memory include Mullainathan et al. (2008) and Schwartzstein and Sunderam (2021).

¹⁹ Related evidence comes from the studies of non-linear pricing. Grubb (2009, 2015) and Ater and Landsman (2013) show that consumers on average choose cell-phone plans and retail banking services with usage allowances that are too large. They also show that consumers who experience overage fees are more likely to switch to plans with larger allowances and end up paying more on average. A similar logic can explain why quality add-ons, which create positive quality surprises, are so profitable (Ellison & Ellison 2005).

attention helps explain consumers' seemingly inconsistent risk attitudes between insurance and investments, which are typically explained via overweighting of unlikely payoffs, and for decoy effects, which are typically explained via menu-dependent preferences. Surprise-driven bottom up attention helps explain time varying risk attitudes in financial markets, and time varying price sensitivity of consumers, which is typically explained through reference point effects. Prominence-driven bottom up attention helps explain framing effects, including many marketing practices, which are typically explained using frame dependent preferences or reference points.

The focus on bottom-up attention does not only offer a psychologically realistic way to unify different forms of choice instability. Relative to the exotic preferences explanations or the optimal attention allocation theories, it yields new predictions grounded in the psychology of bottom up attention and memory. The implications for the welfare effect of reminders is one example, but we discussed many others along the way.

One important avenue for future work is to leverage the explanatory power offered by psychological foundations and explicitly test salience models, in part by employing new measurements. A growing body of applied work uses the term salience to explain the outsized role of some information on decisions. But as we saw, salience is much more than that, because it covers several precise mechanisms operating in different conditions and with different choice implications. To assess these mechanisms, a researcher needs not just data on information and choices, but also measures of contrast, surprise, and prominence, as well as of sensitivity of choices to these measures.

Measuring contrast requires observing at least some attributes of different choice options (e.g. risk, return, and fees for financial assets, price and quality for consumer goods). Measuring surprise additionally requires information on consumers' histories, ideally

longitudinal data on their choices, which also sheds light on memory effects. As we argued in Section 6, such data may help distinguish welfare improving and welfare reducing reminders.

Measuring prominence requires data on what the consumer sees, and ideally what features are most prominent. This is a tough requirement but it is becoming possible in more and more datasets, including in online settings. If a researcher measures what the consumer sees, neuroscientific models of visual attention can be used to determine which locations in the consumer's visual field are most prominent. In a notable example, Li and Camerer (2020) find that goods in locations predicted to attract gaze are more likely to be chosen, even if they are inferior, consistent with the possibility of bottom up attention interfering with goals.

The salience approach to behavior also requires theoretical progress. One obvious open problem is how to think of contrast, surprise and prominence as evolutionary efficient mechanisms. Future theoretical analysis should also seek to assess the role of bottom up attention in phenomena that we did not discuss here but that are particularly relevant in natural choice environments. Here we list a few.

Kahneman (2011) stresses the importance of multiple reference points. When a consumer chooses among many goods at a supermarket, how do the goods in front of him and memories of goods from the past compete in shaping payoff contrast? Answering this question amounts to fully integrating the mechanisms of contrast and surprise.

Relatedly, in our setup all available options are considered, but this is not the case in reality. People frame choices narrowly, by neglecting opportunity costs of a purchase and more generally engaging in mental accounting (Thaler 2008). The evidence from Li and Camerer (2020) as well as research on the attention-DDM (Shimojo et al. 2003, Krajbich & Rangel 2011, Mormann and Russo 2021) shows that driving consumer's attention to one good tends to encourage its selection, perhaps because its prominence interferes with holding alternatives in working memory. To study these issues, one needs a model in which

prominence and memory effects shape a consumer's consideration set. One obvious application of such a model is to habit formation and persistence.

Bottom-up attention might also shed light on other domains of choice. A growing literature suggests that social preferences, fairness judgments, and strategic interactions are also sensitive to seemingly irrelevant changes in framing (Kahneman, Knetsch and Thaler 1985, Krupka Weber 2013, Dufwenberg et al 2011). Such instability might be traced to frames drawing attention to some consequences over others. Or consider attitudes toward ambiguity (Ellsberg 1961). In the absence of objective probabilities, what guides our attention to particular payoffs and their likelihood? How does payoff contrast interact with the construction of scenarios from memory? We would love to know.

Once memory and bottom up attention are taken into account, many choice anomalies begin to make sense. With this progress, we feel that a general understanding of mechanisms behind choice instability is within reach. Reaching this understanding is critical for economic analysis as well as for policy, as our discussion on reminders illustrates.

References

- Abito JM, Salant Y. 2019. The effect of product misperception on economic outcomes: evidence from the extended warranty market. *Rev. Econ. Stud.* 86(6): 2285–2318
- Akerlof G. 1991. Procrastination and obedience. *Am. Econ. Rev.* 81(2): 1-19
- Allais M. 1953. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine. *Econometrica* 21(4): 503–546
- Ariely D, Wallsten TS. 1995. Seeking subjective dominance in multidimensional space: an explanation of the asymmetric dominance effect. *Organ. Behav. Hum. Decis. Process* 63(3): 223-232
- Ater I, Landsman V. 2013. Do customers learn from experience? Evidence from retail banking. *Manag. Sci.* 59(9): 2019-2035
- Barberis N, Huang M. 2008. Stocks as lotteries: the implications of probability weighting for security prices. *Am. Econ. Rev.* 98 (5): 2066-2100
- Barseghyan L, Molinari F, O'Donoghue T, Teitelbaum J. 2013. The nature of risk preferences: evidence from insurance choices. *Am. Econ. Rev.* 103(6): 2499-2529
- Bernheim D, Rangel A. 2009. Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *Q. J. Econ.* 124(1): 51-104
- Bhargava S, Fisman R. 2014. Contrast effects in sequential decisions: evidence from speed dating. *Rev. Econ. Stat.* 96(3): 444-457
- Bodner R, Prelec D. 1994. *The centroid model of context dependent choice*. Manuscript, MIT.
- Bordalo P, Gennaioli N, Shleifer A. 2012. Saliency theory of choice under risk. *Q. J. Econ.* 127(3): 1243-1285
- Bordalo P, Gennaioli N, Shleifer A. 2013a. Saliency and consumer choice. *J. Political Econ.* 121(5): 803-843
- Bordalo P, Gennaioli N, Shleifer A. 2013b. Saliency and asset prices. *Am. Econ. Rev P&P* 103(3): 623-628.
- Bordalo P, Gennaioli N, Shleifer A. 2016. Competition for attention. *Rev. Econ. Stud.* 83(2): 481–513
- Bordalo P, Gennaioli N, Shleifer A. 2019. Memory and reference prices: an application to rental choice. *Am. Econ. Rev P&P* 109 (5): 572–576
- Bordalo P, Gennaioli N, Shleifer A. 2020. Memory, attention and choice. *Q. J. Econ.* 135(3): 1399–1442

- Bose D, Cordes H, Nolte S, Schneider J, Camerer C. 2020. *Decision weights for experimental asset prices based on visual salience*. Working paper
- Broadbent DE. 1958. *Perception and Communication*. New York: Pergamon Press.
- Bruhin A, Manai M, Pinto L. 2020. *Risk and rationality: the relative importance of probability weighting and choice set dependence*. Working paper
- Bushong B, Rabin M, Schwartzstein J. 2021. A model of relative thinking. *Rev. Econ. Stud.* 88(1): 162–191
- Busse M, Pope D, Pope J, Silva-Risso J. 2015. The psychological effect of weather on car purchases. *Q. J. Econ.* 130(1): 371–414
- Calzolari G, Nardotto M. 2017. Effective reminders. *Manag. Sci.* 63(9): 2915-2932
- Célérier C, Vallée B. 2017. Catering to investors through security design: headline rate and complexity. *Q. J. Econ.* 132(3): 1469–1508
- Chang T, Huang W, Wang Y. 2018. Something in the air: pollution and the demand for health insurance. *Rev. Econ. Stud.* 85(3): 1609-1634
- Chetty R, Looney A, Kroft K. 2009. Salience and taxation: theory and evidence. *Am. Econ. Rev.* 99(4): 1145-77
- Chiappori PA, Salanié B, Salanié F, Gandhi A. 2019. From aggregate betting data to individual risk preferences. *Econometrica* 87(1): 1-36
- Clark A. 2013. Whatever next? Predictive brains, situated agents and the future of cognitive science. *Behav. Brain Sci.* 36(3): 181-204
- Conlin M, O'Donoghue T, Vogelsang T. 2007. Projection bias in catalog orders. *Am. Econ. Rev.* 97(4): 1217-1249
- Dertwinkel-Kalt M, Köhler K, Lange MRJ, Wenzel T. 2017. Demand shifts due to salience effects: experimental evidence. *J. Eur. Econ. Assoc.* 15 (3): 626–653
- Dertwinkel-Kalt M, Gerhardt H, Riener G, Schwerter F, Strang L. 2021. Concentration bias in intertemporal choice. *Rev. Econ. Stud.* forthcoming.
- Dertwinkel-Kalt M, Köster M. 2020. Salience and skewness preferences. *J. Eur. Econ. Assoc.* 18(5): 2057-2107
- Desimone R, Duncan J. 1995. Neural mechanisms and selective visual attention. *Annu. Rev. Neurosci.* 18:193-222
- Dessaint O, Matray A. 2017. Do managers overreact to salient risks? Evidence from hurricane strikes. *J. Financ. Econ.* 126(1): 97-121
- Deutsch JA, Deutsch D. 1963. Attention: some theoretical considerations. *Psychol. Rev.* 87: 272– 300

- Dufwenberg M, Gächter S, Hennig-Schmidt H. 2011. The framing of games and the psychology of play. *Games Econ Behav.* 73(2), 459–478
- Dunn K, Bremner JG. 2016. Investigating looking and social looking measures as an index of infant violation of expectation. *Dev. Sci.* 20(6): e12452
- Ellis A, Masatlioglu Y. 2021. Choice with endogenous categorization. *Rev. Econ. Stud.*, forthcoming
- Ellison G, Fischer Ellison S. 2005. Lessons about markets from the internet. *J. Econ. Perspect.* 19(2): 139-158
- Ellsberg D. 1961. Risk, ambiguity, and the Savage axioms. *Q. J. Econ.* 75(4): 643-669
- Enke B. 2020. What you see is all there is, *Q. J. Econ.* 135(3): 1363–1398
- Finkelstein A. 2009. E-z tax: tax salience and tax rates. *Q. J. Econ.* 124(3): 969–1010
- Fishburn P. 1989. Non-transitive measurable utility for decision under uncertainty. *J. Math. Econ.* 18: 187-207
- Frederick S, Novemsky N, Wang J, Dhar R, Nowlis S. 2009. Opportunity cost neglect. *J. Consum. Res.* 36(4): 553-561
- Friedman M, Savage LJ. 1948. The utility analysis of choices involving risk. *J. Political Econ.* 56: 279–304
- Frydman C, Mormann M. 2020. *The role of salience in choice under risk: an experimental investigation.* Working paper
- Gabaix X. 2014. A sparsity-based model of bounded rationality. *Q. J. Econ.* 129(4): 1661–1710
- Gabaix X, Laibson D. 2005. *Bounded rationality and directed cognition.* Working Paper.
- Gagnon-Bartsch T, Rabin M, Schwartzstein J. 2018. *Channeled attention and stable errors.* Working Paper
- Goldin J, Homonoff T. 2013. Smoke gets in your eyes: cigarette tax salience and regressivity. *Am. Econ. J.: Econ Policy* 5(1): 302-336
- Grether DM, Plott CR. 1979. Economic theory of choice and the preference reversal phenomenon. *Am. Econ. Rev.* 69(4): 623-638
- Grubb MD. 2009. Selling to overconfident consumers. *Am. Econ. Rev.* 99(5): 1770-1807
- Grubb MD. 2015. Failing to choose the best price: theory, evidence, and policy. *Rev. Ind. Organ.* 47: 303–340
- Handel B, Schwartzstein J. 2018. Frictions or mental gaps: what's behind the information we (don't) use and when do we care? *J. Econ. Perspect.* 32 (1): 155-78

- Hanna R, Mullainathan S, Schwartzstein J. 2014. Learning through noticing: theory and evidence from a field experiment. *Q. J. Econ.* 129(3): 1311–1353
- Hartzmark SM, Shue K. 2018. A tough act to follow: contrast effects in financial markets. *J. Finance* 73: 1567-1613
- Hastings J, Shapiro J. 2013. Fungibility and consumer choice: evidence from commodity price shocks. *Q. J. Econ.* 128: 1449–1498
- Hawkins J, Blakeslee S. 2004. *On intelligence*. Macmillan.
- Heath TB, Chatterjee S. 1995. Asymmetric decoy effects on lower-quality versus higher-quality brands: meta-analytic and experimental evidence. *J. Consum. Res.* 22(3): 268-284
- Huber J, Payne JW, Puto C. 1982. Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *J. Consum. Res.* 9(1): 90–98
- Itti L, Baldi P. 2009. Bayesian surprise attracts human attention. *Vis. Res.* 49(10): 1295-1306
- Itti L, Koch C. 2001. Computational modelling of visual attention. *Nat. Rev. Neurosci.* 2: 1-11.
- Itti L, Koch C, Niebur E. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 20(11): 1254-1259
- James W. 1890. *The Principles of Psychology*. New York: Henry Holt.
- Kahneman D. 2011. *Thinking, fast and slow*. Macmillan.
- Kahneman D, Chajczyk D. 1983. Tests of the automaticity of reading: dilution of stroop effects by color-irrelevant stimuli. *J. Exp. Psychol.* 9(4): 497-509
- Kahneman D, Knetsch J, Thaler R. 1986. Fairness as a constraint on profit seeking: entitlements in the market. *Am Econ Rev.* 76 (4), 728-741
- Kahneman D, Miller DT. 1986. Norm theory: comparing reality to its alternatives. *Psychol. Rev.* 93(2): 136–153
- Kahneman D, Tversky A. 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47(2): 263–292
- Kahneman D, Tversky A. 1984. Choices, values, and frames. *Am. Psychol.* 39(4): 341–350
- Khaw MW, Li Z, Woodford M. 2021. Cognitive imprecision and small-stakes risk aversion. *Rev. Econ. Stud.* 88(4): 1979–2013
- Karlan D, Ratan AL, Zinman J. 2014. Savings by and for the poor: a research review and agenda. *Rev. Income Wealth* 60: 36-78
- Kontek K. 2016. A critical note on salience theory under risk. *Econ. Lett.* 149: 168-171.
- Koschate-Fischer N, Wüllner K. 2017. New developments in behavioral pricing research. *Rev. Income Wealth* 87(6): 809-875

- Köster M. 2021. *A Multivariate Saliency theory of Choice Under Risk*. Working paper
- Kőszegi B, Rabin M. 2006. A model of reference-dependent preferences. *Q. J. Econ.* 121(4): 1133-1165
- Kőszegi B, Szeidl A. 2013. A model of focusing in economic choice, *Q. J. Econ.* 128(1): 53–104
- Krajbich I, Armel C, Rangel A. 2010. Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* 13(10): 1292-1298
- Krajbich I, Rangel A. 2011. Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc. Natl. Acad. Sci.* 108(33): 13852-13857
- Krupka E, Weber R. 2013. Identifying social norms using coordination games: why does dictator game sharing vary? *J. Eur. Econ. Assoc.* 11(3): 495-524
- Kunreuther H, Ginsberg R, Miller L, Sagi P, Slovic P, Borokan B, Katz N. 1978. *Disaster insurance protection: Public policy lessons*. New York: Wiley
- Lanzani G. 2021. Correlation made simple: applications to saliency and regret theory. *Q. J. Econ.* Forthcoming.
- Lehmann D, Pan Y. 1994. Context effects, new brand entry, and consideration sets. *J. Mark. Res.* 31: 364–374
- Lewandowsky S, Jetter M, Ecker UKH. 2020. Using the president’s tweets to understand political diversion in the age of social media. *Nat. Commun.* 11(5764)
- Li X, Camerer C. 2021. Predictable effects of bottom-up visual saliency in experimental decisions and games. *Q.J. Econ.* Forthcoming.
- Lian C, Ma Y, Wang C. 2019. Low interest rates and risk-taking: evidence from individual investment decisions. *Rev. Financ. Stud.* 32(6): 2107–2148
- Lieder F, Hsu M, Griffiths T. 2014. The high availability of extreme events serves resource-rational decision-making. *Proc. 36th Annu. Conf. Cogn. Sci. Soc.*
- Lieder F, Griffiths T. 2020. Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behav. Brain Sci.* 43, 1-60.
- Lichtenstein S, Slovic P. 1971. Reversals of preference between bids and choices in gambling decisions. *J. Exp. Psychol.* 89(1): 46–55
- Liesefeld HR, Liesefeld AM, Müller HJ, Rangelov D. 2017. Saliency maps for finding changes in visual scenes. *Atten. Percept. Psychophys.* 79: 2190–2201
- Lockwood B, Allcott H, Taubinsky D, Sial A. 2021. *The optimal design of state-run lotteries*. working paper

- Loewenstein G, O'Donoghue T, Rabin M. 2003. Projection bias in predicting future utility. *Q. J. Econ.* 118 (4): 1209-1248.
- Loomes G, Sugden R. 1982. Regret theory: an alternative theory of rational choice under uncertainty. *Econ. J.* 92(368): 805–824.
- Malmendier U, Nagel S. 2011. Depression babies: do macroeconomic experiences affect risk taking? *Q. J. Econ.* 126(1): 373–416
- Mormann M, Russo E. 2021. Does attention increase the value of choice alternatives? *Trends Cogn. Sci.* 25(4): 305 – 315
- Mullainathan S, Schwartzstein J, Congdon WJ. 2012. A reduced-form approach to behavioral public finance, *Annu. Rev. Econ.* 4(1): 511-540
- Mullainathan S, Schwartzstein J, Shleifer A. 2008. Coarse thinking and persuasion. *Q. J. Econ.* 123(2): 577–619
- Nobre A, Mesulam M. 2014. Large scale networks for attentional biases. *The Oxford Handbook of Attention*, A. Nobre and S. Kastner eds.
- Ortmeyer GK, Quelch JA, Salmon WJ. 1991. Restoring credibility to retail pricing. *MIT Sloan Manag. Rev.* 33(1): 55–66
- Paschur T, Schulte-Mecklenbeck M, Murphy R, Hertwig R. 2018. Prospect theory reflects selective allocation of attention, *J. Exp. Psychol. Gen.* 147 (2): 147 – 169.
- Pashler H. 1988. Familiarity and visual change detection. *Percept. Psychophys.* 44, 369-378
- Pashler, H. 1998. *The Psychology of Attention*, MIT Press
- Patrick K, Raab F, Adams M, Dillon L, Zabinski M, Rock C, Griswold W, Norman G. 2009. A text message-based intervention for weight loss: randomized controlled trial. *J. Med. Internet. Res.* 11(1):e1
- Peters RJ, Itti, L. 2007. Beyond bottom-up: incorporating task-dependent influences into a computational model of spatial attention. In: *2007 IEEE Conf. computer vision and pattern recognition*: 1-8
- Pratt JW, Wise DA, Zeckhauser R. 1979. Price differences in almost competitive markets, *Q. J. Econ.* 93(2): 189–211
- Radbruch J, Schiprowski A. 2021. *Interview sequences and the formation of subjective assessments*. Working paper
- Rao R, Ballard D. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2: 79–87
- Remington RW, Johnston JC, Yanis S. 1992. Involuntary attentional capture by abrupt onsets. *Percept. Psychophys.* 51: 726-744.

- Rubinstein A. 1988. Similarity and decision-making under risk. *J. Econ. Theory* 46: 145–153
- Savage LJ. 1954. *The Foundations of Statistics*. John Wiley & Sons
- Schwartzstein J. 2014. Selective attention and learning. *J. Eur. Econ. Assoc.* 6(1): 1423–1452
- Schwartzstein J, Sunderam A. 2021. Using models to persuade. *Am. Econ. Rev.* 111(1): 276–323
- Shiffrin RM, Gardner GT. 1972. Visual processing capacity and attentional control. *J. Exp. Psychol.* 93(1):72-82
- Shimojo S, Simion C, Shimojo E, Scheier C. 2003. Gaze bias both reflects and influences preference. *Nat. Neurosci.* 6(12): 1317-22
- Schulte-Mecklenbeck M, Johnson J, Böckenholt U, Goldstein D, Russo J, Sullivan N, Willemssen M. 2017. Process-tracing methods in decision making: on growing up in the 70s. *Curr Dir Psychol Sci.* 26(5): 442-450.
- Simons DJ, Chabris CF. 1999. Gorillas in our midst: sustained inattention blindness for dynamic events. *Perception* 28(9): 1059-1074
- Simonsohn U. 2006. New Yorkers commute more everywhere: contrast effect in the field. *Rev. Econ. Stat.* 88(1): 1-9
- Simonsohn U, Loewenstein G. 2006. Mistake #37: the effect of previously encountered prices on current housing demand, *Econ. J.* 116(508): 175-199
- Simonson I. 1989. Choice based on reasons: the case of attraction and compromise effects. *J. Consum. Res.* 16(2): 158–174
- Simonson I, Tversky A. 1992. Choice in context: tradeoff contrast and extremeness aversion. *J. Mark. Res.* 29(3): 281–295
- Sims CA. 2003. Implications of rational inattention. *J. Monet. Econ.* 50(3): 665-690
- Soltani A, De Martino B, Camerer C. 2012. A range normalization model of context dependent choice: a new model and evidence. *PLOS Comput. Biol.* 8(7): e1002607
- Sydnor J. 2010. (Over)insuring modest risks. *Am. Econ. J.: Appl. Econ.* 2(4): 177-199
- Stango V, Zinman J. 2014. Limited and varying consumer attention: evidence from shocks to the salience of bank overdraft fees. *Rev. Financ. Stud.* 27(4): 990–1030
- Sussman AB, Alter AL. 2012. The exception is the rule: underestimating and overspending on exceptional expenses. *J. Consum. Res.* 39(4): 800–814
- Taylor S, Thompson S. 1982. Stalking the elusive vividness effect. *Psychol. Rev.* 89(2): 155–181
- Thaler R. 1985. Mental accounting and consumer choice. *Mark. Sci.* 4(3): 199-214

- Thaler R, Sunstein CR. 2008. *Nudge: improving decisions about health, wealth and happiness*. Yale University Press
- Treisman A, Gelade G. 1980. A feature integration theory of attention. *Cogn. Psychol.* 12(1): 97-136
- Treisman A, Gormican S. 1988. Feature analysis in early vision: evidence from search asymmetries. *Psychol. Rev.* 95(1): 15-48
- Tversky A, Kahneman D. 1991. Loss aversion in riskless choice: a reference-dependent model. *Q. J. Econ.* 106(4): 1039-1061
- Tversky A, Kahneman D. 1992. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* 5: 297–323
- Tversky A, Slovic P, Kahneman D. 1990. The causes of preference reversals. *Am. Econ. Rev.* 80(1): 204–217
- Tversky A, Simonson I. 1993. Context-dependent preferences. *Manag. Sci.* 39(10): 1179-1189
- Woodford M. 2012. Prospect theory as efficient perceptual distortion. *Am. Econ. Rev P&P* 102(3): 41-46
- Woodford M. 2020. Modeling imprecision in perception, valuation, and choice. *Annu. Rev. Econ.* 12(1): 579-601
- Wu C, Cosguner K. 2020. Profiting from the decoy effect: a case study of an online diamond retailer. *Mark. Sci.* 39(5): 974-995