Attention, Strategy, and the Human Mind

Thesis by
Xiaomin Li

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2020
Defended September 30, 2020

Xiaomin Li
ORCID: 0000-0002-1286-4012

ACKNOWLEDGEMENTS

My first conversation with Colin Camerer happened before he even remembers. It was at a conference dinner in Miami at the Society of Neuroeconomics, where he told me that the Ph.D. program at Caltech was unique but challenging. And it proved to be exactly what he first said. I am deeply grateful having gone through this extraordinary journey in my life with Colin's guidance and advice. I know how precious it was for Colin to spend time and effort talking to me about every detail of my research and my life statuses like my visa application or emotional feelings.

I am also thankful to Ralph Adolphs, who provided me with all the lab resources when I first started a project. I remember that Ralph helped me with polishing my first, and probably the most terrible draft ever: my second-year paper. I cannot imagine how much time he put in it, and got nothing in return for that. I am thankful to Lawrence Jin, my committee member and my mentor outside of neuroeconomics, who generously spent hours listening to my project, which was unrelated to his field, and gave me numerous career advice. Also to my committee chair Omer Tamuz, who kindly took on the responsibility to instruct a student who does nothing about economic theory. Thanks also for being our best game theory teacher and trusting me to be his TA in the later years.

I would also like to thank Antonio Rangel, who listened to many of my crazy ideas and organized the best seminar class ever on neuroeconomics. I enjoy going to this class every week and just listen to his sharp thoughts on various researches. Also thanks to John O'Doherty for being our director of the program and who shared with me precious career advises.

I am also thankful to the entire HSS department, where I felt a sense of belonging through every event that has been organized. I would specially like to thank Laurel Auchampaugh, Barbara Estrada, and Kapauhi Stibbard, without whom my life would have been much harder having to bother with lots of administrative affairs.

I would like to thank everyone in Camerer's group, including all the visiting scholars and summer students, for all the discussions we had every Wednesday. It is hard to count how much I learned from these meetings. I especially thank Taisuke Imai, to whom I was assigned as a mentee and who treated me with kindness and patience when I just started. I also need to say thank you to Nina Solovyeva, who was a mentee of mine, but who helped me out with the experiment in chapter four.

I also want to say thank you to Adam Brandenburger, without whom I would not have even started or thought of pursuing a Ph.D. degree. Adam is my lifetime mentor and friend, who kept being supportive every time I feel hopeless and puzzled.

I am also thankful to my parents, who provided me with the happiest family so that I am never afraid to challenge anything in the world including researches.

I am especially thankful to Yi Xu, a computer vision scientist, the smartest man I ever know, who has listened to every single piece of my research and always answered the hardest and the most important question: “Does this make any sense to you?” This sometimes refered to an idea; sometimes, it could be a piece of result that I just got or a thought from another paper. He helped me defeat the deepest fear and the largest enemy along the way, which is that my research is meaningless to anyone other than myself. On top of that, he also provided me with a large amount financial support to continue chasing my academic dreams.

ABSTRACT

The current dissertation chapters try to discover the role of visual attention in decision making from three different perspectives: 1) how attention bias affects strategic decision makings, 2) how to model eye movement data to better understand strategic decisions, 3) how to manipulate simple choices through visual saliency.

The second chapter introduces a series of novel image games, where players need to match, hide, or seek against other players. We apply a pure computational way: the state-of-art visual saliency algorithm, Saliency Attentive Map (SAM) to measure visual saliency. We find that visual saliency can predict strategic behaviors well. The concentration of salience is correlated with the rate of matching when players are both trying to match location choices (r=.64). In hider-seeker games, all players choose salient locations more often than predicted in equilibrium, creating a "seeker's advantage" (seekers win 9% of games rather than the 7% predicted in equilibrium). The 9% win rate is robust for paying higher stakes and using a between-subjects design. Salience-choice relations are consistent with cognitive hierarchy and level-k models in which strategically naive level 0's are biased toward salience, and higher-level types are not directly biased toward salience, but choose salient locations because they believe lower-level types do. Other links between salience as understood in psychology and hypothesized in economics are suggested.

The third chapter is a continuation of the second chapter, but with a different emphasis. The third chapter proposes a way to dynamically model gaze transitional data in games utilizing a class of machine learning model: hidden markov models(HMM). The HMM model reveals how the attentional bias affects strategies on different time point. Besides, this model well connects to the k level behavioral method and can make novel predictions on strategic levels. With further containing the fixation duration data, we developed a continuous-time hidden Markov model (cgtHMM), which can be used to predict how exactly time pressure changes choices and the seeker's advantage.

Distinct from the other two, chapter four aims at manipulating binary choice outcomes through the change of visual saliency distribution under SAM. We design a value-based choice paradigm where both the reward property and the attention property are well separated and controlled. The experimental results indicate that visual saliency can enhance the choice correction rates when the more rewarding outcome

is also labeled salient. It can also shorten the decision time needed. Such a result can be explained by a saliency-enhanced rational inattention model by incorporating attention factors in the traditional RI model.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Chapter 1

INTRODUCTION

1.1 General introduction

Challenges of the classic economics theory originated as early as the theory of bounded rationality (Simon, 1956; Simon, 1997), in which a decision maker, instead of fully optimizing, is just seeking a satisfactory solution under limited cognition. What followed afterward is the emergence of behavioral economics. Researchers try to raise particular questions on each concept in economics. Camerer (1999) concludes four classes of economic principles that need to be reconsidered: expected utility theory, exponential discounting, social utilities, and equilibrium. The last two decades were a golden age for behavioral economics 1.0, during which immense progress has been made towards these four directions, both from a theoretical perspective and an empirical perspective. For example, a large body of literature on prospect theory (Tversky and Kahneman, 1992) provides an alternative method to the expected utility theory. More diverse discounting functions, including the famous hyperbolic one (Laibson, 1997; Frederick, Loewenstein, and O'donoghue, 2002), have gained more popularity now than the traditional exponential discounting way. The concept of empathy or fairness is becoming increasingly essential to reshape our understanding of what the actual utility is (Fehr and Schmidt, 1999; Fehr and Gachter, 2000). The equilibrium is not the only status in the economic system that we care about anymore. And in fact, much evidence has shown that we are further away from the equilibrium most of the time in reality, and arriving there is never easy (Camerer and Hua Ho, 1999; Camerer, Ho, and Chong, 2002; Camerer, 2011).

Most of these studies begin with questioning one kind of hypothesis based on intuition or heuristics. Specific experiments are conducted to provide supporting evidence to the doubts and, most of the time, results are indirectly inferred through choice data. Most researches will assume a psychological perturbed model and estimate the corresponding psychological effects using the choice data. Such a post-hoc methodology works to make things progress, but probably it is not the most efficient way nor a well-founded way. The reason of only using choice data comes from the lack of measurements in different mental activities. In this aspect, psychologists and cognitive scientists move way ahead of us. They possess many

treasured techniques to measure mental activities related to economic decisions like eye-tracking or other psychophysics tools.

Indeed, the progress of behavioral economics is also the reunion process of psychology and economics (Camerer, 1999). But most of the time, economists are borrowing only the concepts and results from psychology, for example, emotion, attention, or memory, and leave behind the methodology. We see very few studies in economics which use eye-tracking techniques to study attention or use the psychophysics monitoring package to study emotion.[1] One possible reason for not doing this might be that economists are afraid that a concrete definition will make the concept lose its generalization power. Of course, leaving a concept always as general as possible to be only one word or one sentence in theory will make it omnipotent. For instance, one can make a perfect definition of attention as the distribution of cognitive resources. While on the other hand, if we never concretize such a "perfect definition", this concept will be useless because no one knows how to observe and measure cognitive resources. No one knows what the distribution of it is.

What this thesis did was to concretize the concept of attention and present several application cases of economics on the measured attention. We limit the discussion of attention in this dissertation only within the field of visual attention. Doing so certainly makes this work's scope narrow, as many of the attention topics outside of the visual world, for example, in the auditory system or on the abstract level. Though studying the visual attention role itself is important enough as we humans are "visual animals" and more than 50 percent of the cortex, the surface of the brain, is devoted to processing visual information (Kandel et al., 2000).

Extensive work on psychology and neuroscience contributes to the knowledge of the system of visual attention. They are trying to answer a question: what makes an item attentive. We will introduce them in the following sections.

1.2 The concept of salience

One important ability for animals is to prioritize some of the information stimuli in the environment selectively. This process starts very early in the stage, and forms a so-called "saliency map". Saliency map topographically encodes stimulus conspicuity is formed on several brain areas, including superior colliculus and primary visual cortex (White et al., 2017; Veale, Hafed, and Yoshida, 2017). This

[1] Including skin conductance, heart rate, and face muscle changes.

"saliency map" encodes an order of importance of the content on the entire visual scene coming from our retina. The entire computation process of the saliency map is complicated. Many pathways and brain areas are involved (Veale, Hafed, and Yoshida, 2017). There are two types of major mechanisms involved to determine the order of priority. One type is feature-driven, sometimes is also called the "bottom-up" process. This type purely computes a one-dimensional map on each point on the visual scene using purely low level features including brightness, color, oriented edge and motion (White et al., 2017). Different low level features compete within themselves and then the most salient location comes out. This process happens first, followed by the second type of process: the goal-directed process ("top-down"), which integrates the outcome of the feature-driven mechanism and the information of tasks and goals.

Both types of mechanisms decide the saccades and micro-saccades behaviors, and therefore the eye fixation maps.

1.3 Bottom-up and top-down salience

In psychology, attention-getting features that are salient are usefully divided into two categories: Bottom-up and top-down (Frintrop, Rome, and Christensen, 2010). Bottom-up features are based on sensory physical features such as color, contrast, and orientation. Bottom-up feature salience is devoid of personal meaning and independent of task goals. To the extent that human visual systems work the same way all around the world, the attention-grabbing salience of bottom-up features should be universal. As a result, attentional data from many diverse subject pools can establish what is bottom-up. For example, the stick figure "I" is perceived, bottom-up, as a black vertical line of a certain length, with slightly extended top and bottom horizontal lines on top of the vertical line, surrounded by contrast with a white background. Primates, babies, and adult humans all have the same immediate perception. From the bottom-up view the stimulus "I" is not made more or less salient because it is a capital letter, a pronoun for "myself" in English, the Roman numeral one, or reminds you of an iconic column your interior decorator suggested for the front porch of your house. All the latter forms of salience are "top-down": they use semantic knowledge about the (local, acculturated) world, include what is personally relevant, valued, familiar, and novel. Bottom-up attention seems to emerge from rapid visual neural processing which is essentially common to humans (excepting disorders) and many other species. Inspired by a deep understanding of visual cortex, a series of

progressively improving algorithms have been developed to use visual images as inputs, and output predictions about where people will look in the first 1000msec of processing (Harel, Koch, and Perona 2007; Itti, Koch, and Niebur 1998; Judd et al. 2009). However, as the "I" example suggests, many types of top-down meaning and value have been shown to be salient because they also grab attention. Variables that increase top-down salience in perception and choice include prior beliefs, recent choice history (Awh, Belopolsky, and Theeuwes, 2012), familiarity and novelty (Itti and Baldi, 2009), value for consumer goods (Towal, Mormann, and Koch, 2013), and self-reported "meaning" (Henderson and Hayes, 2017). Evidence also suggests that top-down attention allocation occurs more slowly than bottom-up attention. The SAM algorithm we will use is tuned using human free gaze data. Since the humans have no special goal to achieve during free viewing, the salience that is being captured algorithmically is probably a combination of bottom-up features, and top-down features (such as familiarity, novelty, special attention to faces, etc.) which grab the attention of the particular subject population.

1.4 A computer vision approach for visual saliency modeling

Itti-Koch Saliency Map

Inspired by the attention control function from the primate visual systems, vision scientists take the reverse engineering approach and try to build models that can output a similar result as the attentional map generated by the biological visual system.

The first well-known model is the so-called "Itti-Koch bottom-up saliency map" (Itti, Koch, and Niebur, 1998), which takes thoughts from Koch and Ullman (1987), and uses a purely bottom-up approach to compute for a single topographical saliency map. Figure 1.1 presents the general model structure of the "Itti-Koch" model.

As we walk through the model, we can see that every single procedure of this model reflects a biological foundation. This model takes an RGB image as its input. The first thing it does is to linear filter the image into three basic dimensions: color (C), intensity (I), and orientation (O), just like the way our visual system also has certain types of specialized neurons from the retina to the primary visual cortex to encode intensity, color, and orientations, individually (Kandel et al., 2000). The critical procedure for saliency computation is to get a contrast measurement under each feature. Such contrast is computed by a set of "center-surround" operations

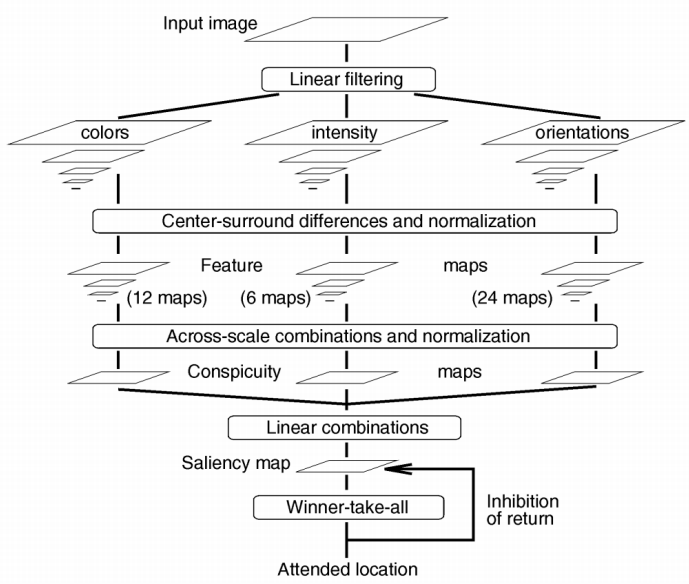

Figure 1.1: Itti-Koch model for bottom-up visual saliency.

The general model structure for the model in Itti, Koch, and Niebur (1998).

just akin to the "center-surround" responses in visual receptive fields.[2] Then, after normalizing such contrast under each of the three dimensions: C, I, O, the algorithm just takes the average of the contrast measure under each dimension as the final saliency map of a given image.

Though this model is so simple to construct and implement compared to many current visual saliency models, it can already partition the stand-out items very well on simple images (e.g., Figures 4 and 5 in (Itti, Koch, and Niebur, 1998)).

Interpretable models before deep neural networks

This computational approach has been later proved to be successful and meaningful. Later studies further tried to improve the model performance to be even closer to how humans look.

But before that, a well-defined measurement of how close a model is to the human gaze pattern becomes important. From Harel, Koch, and Perona (2007) onward, people started to take the experimental approach and get real human gaze data under free viewing tasks for a short duration (3-5s). Fixation maps were then calculated and used as the model target. The performance of a saliency model will

[2]For example, some of the retinal ganglion cells have such center-surround receptive field that some cells will be activated by a bright spot in the middle and dark circle around it.

then be evaluated based on ROC curves drawn between the model prediction and the fixation maps. With these human fixation maps, a pure bottom-up approach is clearly not enough. Many models were then developed to improve at some perspectives, including center biases, top-down semantics, Bayesian priors, and more (Hou and L. Zhang, 2009; Judd, Ehinger, et al., 2009; Harel, Koch, and Perona, 2007; J. Zhang and Sclaroff, 2013).

Up until now, all the models were fully interpretable and crafted 100% by humans. We already learned a lot about what makes a pixel salient, though there is still room for performance to grow.

Standardized metrics and the deep learning approach

Just like other computer vision sub-fields, visual saliency models now also have standardized criteria and data sets for performance checks. Public available data sets like MIT300 (Judd, Durand, and Torralba, 2012) and CAT2000 (Borji and Itti, 2015) contain a large number of fixation maps on more than 4000 images across 200 categories. Besides the uniformed data sets, standardized evaluation metrics including the area under roc curve (AUC), the normalized scanpath saliency (NSS), the correlation coefficient (CC), and others are often applied together now to evaluate the contribution of a new model. The entire industry has a clear goal now: predict how people look. And this goal has been quantified as to reach the golden standard, the human benchmark values under each evaluation metric. [3]

Under the gigantic move of deep neural network models, the state-of-art saliency models (Kummerer, Wallis, Gatys, et al., 2017; Jia and Bruce, 2020) have elevated the AUC performance score to 0.87 where the human benchmark is 0.93, and an ideal perfect model is 1. The ceiling of hand-crafted feature models (non-deep neural network) produces an AUC metric of 0.81 (E. Erdem and A. Erdem, 2013; Harel, Koch, and Perona, 2007; Judd, Ehinger, et al., 2009), which is also reasonably well, while all these models are fully explainable.

1.5 The attentional Drift-Diffusion-Model (aDDM)

Over the past decades, accumulation models have been proved to be successful for modeling attention in decision making. Such accumulation models usually require eye fixation data or mouse trajectory data besides the choice data. The fundamental

[3] This is caused by the individual differences in the fixation patterns. The golden standard can be calculated by calculating each evaluation score using the average fixation density as a model to predict another group of people's gaze.

basis for it comes from the active information collecting process initiated by the decision maker. Such a collecting process is reflected by eye movements or other process form data, and can be used to predict choice by the accumulation models.

The attention drift diffusion model(aDDM) is the most well-known computational model and has been tested in various kinds of tasks including purchasing behavior, perceptual decision making, and so on (Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011; Krajbich, Lu, et al., 2012; Hutcherson, Bushong, and Rangel, 2015; Tavares, Perona, and Rangel, 2017; Kummerer, Wallis, and Bethge, 2018). It is a special application of the general drift diffusion model (Ratcliff and McKoon, 2008; Ratcliff, Smith, et al., 2016) on the realm of perceptual decision making. A standard drift diffusion model describes the dynamic process of evidence accumulation for a binary choice task, and can predict two things: choice outcomes and response time distribution. A key innovation from aDDM is that it assumes a discount factor at each time point for the unfixed item. In that case, the relative decision value (RDV) between two choice alternatives at each time t becomes:

$$V_t = V_{t-1} + d(r_{attended} + \theta r_{notattended} + \epsilon_t)$$

where $r_{attended}$ and $r_{notattended}$ denote for the values of the two alternatives and $\theta \in [0, 1]$ captures the attentional bias towards the item that is currently being fixed at. ϵ_t is the Gaussian white noise term.

The aDDM model is definitely a milestone model for exploring the computational confound of visual attention. Besides that aDDM can make reasonable explanations of RT distribution and choice psychometrics, it further sheds light on the exogenous manipulation directly on the saliency property of choice alternatives (Armel, Beaumel, and Rangel, 2008; Shimojo et al., 2003; Milosavljevic et al., 2012).

Though there have been efforts made to extend the application scenario to contain more than two items, for example, in a trinary problem (Krajbich and Rangel, 2011), the mathematical structure of the DDM model by itself will make it very hard to generalize to more items further. In another word, there is no general solution concept applying to every number of choice alternatives.

1.6 The direction of this dissertation

The computation model of visual saliency is the approach we advocate in this dissertation about "concretizing" attention in economics applications and other decision-making scenarios.

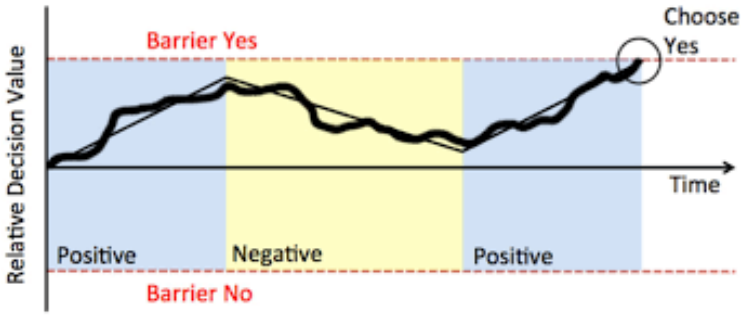

Figure 1.2: The attentional drift diffusion model.

The dynamic change of the relative stimulus value on discrete time for a yes-no question. A choice will be made if one of the barriers is reached.

In chapters two and three, we show that the visual saliency model can powerfully predict the effect of saliency on different game structures quantitatively. Such a hypothesis has been long realized by game theorists and was referred to as the "focal point" effect. With this saliency measurement, we construct a model that is not restricted to any specific game structure and can predict choices better than other competing models. We also did eye-tracking analysis on this experiment ,the result of which consolidates the use of the computational saliency model that the attention measured by the eye tracker is the same as the SAM prediction.

HMM model can further provide us another way of understanding the strategy forming process. In chapter three, we formalize the hidden Markov model in strategic games that eye fixation transitions map to the mental states transitions. Considering the fixation duration time as an input of the model can make it capable of predicting the behaviors under time pressure.

Chapter four provides another application of the computational saliency model in a binary choice decision problem. People generate bias on saliency and make mistakes. We show that the saliency and the valuations contribute together to a decision. We can easily find examples corresponding to this phenomenon in real life. For example, one may take special consideration of the most attentive item placed on the merchant's focal spot by the merchant. She decides to buy it, but finds out later that this is not a valuable item for her.

Overall, this dissertation borrows both the scientific concept of visual saliency and a portable measure from cognitive neuroscience. We concentrated on three economic perspectives where saliency plays an important role: 1) strategic choices, 2) eye fixation data, and 3) simple choice problems. Experiments are later designed and

conducted for each of those three perspectives. Different computational models bridging the novel attention related data set to decisions are finally constructed and estimated to better understand and predict strategic or non-strategic decisions. All models included have the potentials to be generalized and apply to other domains of problems. Limitations of this work and possible future directions are pointed out at the end of this dissertation.

References

Armel, Carrie, Aurelie Beaumel, and Antonio Rangel (2008). "Biasing simple choices by manipulating relative visual attention". In: *Judgment and Decision making* 3(5), pp. 396–403.

Awh, Edward, Artem V. Belopolsky, and Jan Theeuwes (Aug. 2012). "Top-down versus bottom-up attentional control: a failed theoretical dichotomy". In: *Trends in Cognitive Sciences* 16(8), pp. 437–443. ISSN: 1364-6613. DOI: 10.1016/j.tics.2012.06.010. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3426354/ (visited on 05/30/2018).

Borji, Ali and Laurent Itti (2015). "Cat2000: A large scale fixation dataset for boosting saliency research". In: *arXiv preprint arXiv:1505.03581*.

Camerer, Colin (1999). "Behavioral economics: Reunifying psychology and economics". In: *Proceedings of the National Academy of Sciences* 96(19), pp. 10575–10577.

Camerer, Colin (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong (2002). "Sophisticated experience-weighted attraction learning and strategic teaching in repeated games". In: *Journal of Economic theory* 104(1), pp. 137–188.

Camerer, Colin and Teck Hua Ho (1999). "Experience-weighted attraction learning in normal form games". In: *Econometrica* 67(4), pp. 827–874.

Erdem, Erkut and Aykut Erdem (2013). "Visual saliency estimation by nonlinearly integrating features using region covariances". In: *Journal of vision* 13(4), pp. 11–11.

Fehr, Ernst and Simon Gachter (2000). "Cooperation and punishment in public goods experiments". In: *American Economic Review* 90(4), pp. 980–994.

Fehr, Ernst and Klaus M. Schmidt (1999). "A theory of fairness, competition, and cooperation". In: *The Quarterly Journal of Economics* 114(3), pp. 817–868.

Frederick, Shane, George Loewenstein, and Ted O'donoghue (2002). "Time discounting and time preference: A critical review". In: *Journal of Economic Literature* 40(2), pp. 351–401.

Frintrop, Simone, Erich Rome, and Henrik I. Christensen (2010). "Computational visual attention systems and their cognitive foundations: A survey". In: *ACM Transactions on Applied Perception (TAP)* 7(1), p. 6.

Harel, Jonathan, Christof Koch, and Pietro Perona (2007). "Graph-based visual saliency". In: *Advances in Neural Information Processing Systems*, pp. 545–552.

Henderson, John M. and Taylor R. Hayes (2017). "Meaning-based guidance of attention in scenes as revealed by meaning maps". In: *Nature Human Behaviour* 1(10), p. 743.

Hou, Xiaodi and Liqing Zhang (2009). "Dynamic visual attention: Searching for coding length increments". In: *Advances in Neural Information Processing Systems*, pp. 681–688.

Hutcherson, Cendri A, Benjamin Bushong, and Antonio Rangel (2015). "A neurocomputational model of altruistic choice and its implications". In: *Neuron* 87(2), pp. 451–462.

Itti, Laurent and Pierre Baldi (2009). "Bayesian surprise attracts human attention". In: *Vision Research* 49(10), pp. 1295–1306.

Itti, Laurent, Christof Koch, and Ernst Niebur (1998). "A model of saliency-based visual attention for rapid scene analysis". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(11), pp. 1254–1259.

Jia, Sen and Neil DB Bruce (2020). "Eml-net: An expandable multi-layer network for saliency prediction". In: *Image and Vision Computing*, p. 103887.

Judd, Tilke, Frédo Durand, and Antonio Torralba (2012). "A benchmark of computational models of saliency to predict human fixations". In:

Judd, Tilke, Krista Ehinger, et al. (2009). "Learning to predict where humans look". In: *Computer Vision, 2009 IEEE 12th international conference on*. IEEE, pp. 2106–2113. ISBN: 1-4244-4420-9.

Kandel, Eric R et al. (2000). *Principles of Neural Science*. Vol. 4. McGraw-hill New York.

Koch, Christof and Shimon Ullman (1987). "Shifts in selective visual attention: towards the underlying neural circuitry". In: *Matters of Intelligence*. Springer, pp. 115–141.

Krajbich, Ian, Carrie Armel, and Antonio Rangel (2010). "Visual fixations and the computation and comparison of value in simple choice". In: *Nature Neuroscience* 13(10), p. 1292.

Krajbich, Ian, Dingchao Lu, et al. (2012). "The attentional drift-diffusion model extends to simple purchasing decisions". In: *Frontiers in Psychology* 3, p. 193.

Krajbich, Ian and Antonio Rangel (2011). "Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions". In: *Proceedings of the National Academy of Sciences* 108(33), pp. 13852–13857.

Kummerer, Matthias, Thomas SA Wallis, and Matthias Bethge (2018). "Saliency benchmarking made easy: Separating models, maps and metrics". In: *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 770–787.

Kummerer, Matthias, Thomas SA Wallis, Leon A Gatys, et al. (2017). "Understanding low-and high-level contributions to fixation prediction". In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4789–4798.

Laibson, David (1997). "Golden eggs and hyperbolic discounting". In: *The Quarterly Journal of Economics* 112(2), pp. 443–478.

Milosavljevic, Milica et al. (2012). "Relative visual saliency differences induce sizable bias in consumer choice". In: *Journal of Consumer Psychology* 22(1), pp. 67–74.

Ratcliff, Roger and Gail McKoon (2008). "The diffusion decision model: theory and data for two-choice decision tasks". In: *Neural Computation* 20(4), pp. 873–922.

Ratcliff, Roger, Philip L Smith, et al. (2016). "Diffusion decision model: Current issues and history". In: *Trends in Cognitive Sciences* 20(4), pp. 260–281.

Shimojo, Shinsuke et al. (2003). "Gaze bias both reflects and influences preference". In: *Nature neuroscience* 6(12), pp. 1317–1322.

Simon, Herbert (1956). "Rational choice and the structure of the environment." In: *Psychological Review* 63(2), p. 129.

Simon, Herbert (1997). *Models of bounded rationality: Empirically grounded economic reason*. Vol. 3. MIT press.

Tavares, Gabriela, Pietro Perona, and Antonio Rangel (2017). "The attentional drift diffusion model of simple perceptual decision-making". In: *Frontiers in Neuroscience* 11, p. 468.

Towal, R. Blythe, Milica Mormann, and Christof Koch (2013). "Simultaneous modeling of visual saliency and value computation improves predictions of economic choice". In: *Proceedings of the National Academy of Sciences* 110(40), E3858–E3867.

Tversky, Amos and Daniel Kahneman (1992). "Advances in prospect theory: Cumulative representation of uncertainty". In: *Journal of Risk and Uncertainty* 5(4), pp. 297–323.

Veale, Richard, Ziad M Hafed, and Masatoshi Yoshida (2017). "How is visual salience computed in the brain? Insights from behaviour, neurobiology and modelling". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 372(1714), p. 20160113.

White, Brian J et al. (2017). “Superior colliculus encodes visual saliency before the primary visual cortex”. In: *Proceedings of the National Academy of Sciences* 114(35), pp. 9451–9456.

Zhang, Jianming and Stan Sclaroff (2013). “Saliency detection: A boolean map approach”. In: *Proceedings of the IEEE international conference on computer vision*, pp. 153–160.

Chapter 2

SALIENCY IN EXPERIMENTAL GAMES

2.1 Introduction

Salience is whatever makes a feature of a stimulus grabs attention quickly. As behavioral economists explore the implications of scarce attentional resources on choices and market outcomes, a deeper understanding of salience could prove useful in economics. This paper is about one particular aspect of salience: visual salience as a source of focality in games. Pure coordination (matching) and hider-seeker games will be compared.

Schelling (1980) famously noted that how strategies in games are "labeled" mentally – whether by visual salience, language, cultural importance, precedent, or in other "psychologically prominent" ways – seemed to be necessary to predict what people would actually do in certain types of games.

Three decades later Rubinstein (1991) echoed this sentiment, writing (p. 909)

> It is argued that a good model in game theory has to be realistic in the sense that it proves a model for the *perception* of real life social phenomena (emphasis ours)

In games, focality depends on a *shared* concept of salience. Focal strategies must stand out *and* whether people think strategies will stand out to others makes a difference. Visual neuroscience has crafted clear predictions about what most humans notice. Aspects of human perception that are universal – e.g., noticing the color red – are a promising basis for shared salience.

Let us begin with Schelling's (1960) famous map example (see Figure 2.1a). The map shows small square houses, a pond in the lower left, two places marked x and y, and a river running horizontally through the lower third of the map, with a bridge located across the river. Schelling wrote:

> Two people parachute unexpectedly into the area shown, each with a map and knowing the other has one, but neither knowing where the other

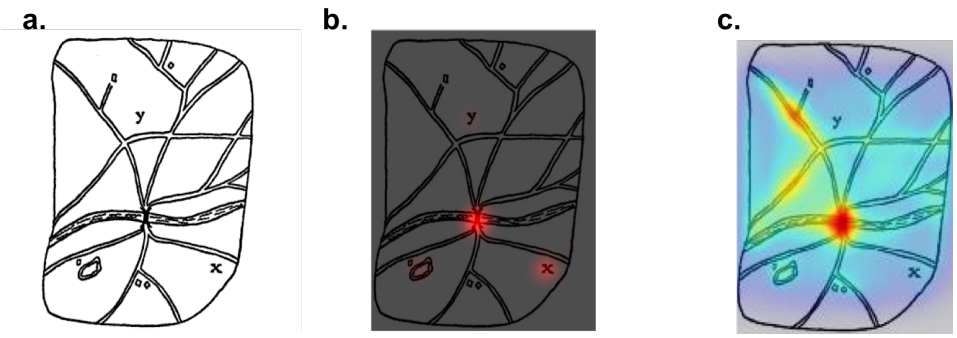

Figure 2.1: Schelling's map revisited.

(a) Original map. (b) Choice frequencies heatmap, the redness indicates choice frequency. (c) SAM algorithm predicted saliency heatmap.

> has dropped nor able to communicate directly. They must get together quickly to be rescued. Can they study their maps and "coordinate" their behavior? (p. 56)

Schelling reported that seven of the eight people that he asked chose to rendezvous at the bridge.

In a proper experiment, N=61 UCLA students, who were incentivized to match, chose the bridge 59% of the time (see Figure 2.1b).[1] Did they choose the bridge because it grabbed their attention? To find out, we fed the map image taken right from Schelling's book into a computational neural network algorithm carefully designed to predict which areas of *any* image are highly salient to average people. (This SAM algorithm is discussed further below).

The resulting SAM map says that the bridge is highly salient (Figure 2.1c). The algorithm predicts that the bridge is salient because it has the most dark/light contrast, and is close to the center of the map. However, note that SAM says that the "x spot" is *not* salient, even though it was chosen by 25% of the subjects. The SAM algorithm, based on what people tend to look at in general, misses some of what makes a map feature salient to people who have the shared goal of trying to coordinate. After presenting our results on games, in the Discussion section at the end of this paper we address how psychological visual salience is related to other concepts of salience used in economics.

[1] These data were collected in conjunction with Milica Moormann and Alec Smith.

What we do that is new

Our study goes further than previous research in three ways

1. Salience is predicted by an underlying neuro-computational theory of which features of an information display people look at first (the SAM algorithm). SAM salience is a good predictor of choices in matching games; it is highly correlated with what image locations players choose, and across games, higher matching rates correlate with more concentrated salience.

2. We estimate salience-enhanced level-k and cognitive hierarchy models and use them to portably explain data in matching games based on estimates of salience strength from hider-seeker games.

3. A dynamic hidden Markov model (HMM) of salient attention over time is constructed from both eyetracking and choice data. Because the model describes the time course of attention to salient locations, it can predict exactly how choices differ under time pressure. It accounts for an increase in the seeker's advantage in a new experimental condition if players are under time pressure. The seeker's advantage increases because the hider does not have time to find an unsalient location to hide at.

Section 2 discusses some classic experimental studies of focality in games. Then in section 3, we introduce the visual image paradigm and algorithms that predict what features are salient. Section 4 describes our experimental design and basic results, and organizes results using a salience-enhanced cognitive hierarchy (SCH) model. We measure how well it can explain data, and whether it can portably make cross-game predictions.

The last section concludes by discussing salience more generally in recent game theory and economic applications, and pointing to future research.

2.2 Background on focality in games

There was a long lag between Schelling's early discussion and later bursts of careful experimentation on focality.

Bacharach (1993); Bacharach and Bernasconi (1997) proposed several principles underlying focality in matching choices from sets of objects. Their idea was that people naturally categorized objects into subsets and chose from distinctive objects

a. **b.**

Figure 2.2: Actual photograph of a deworming bracelet and associated SAM saliency map.

Original image of deworming bracelet (left) and saliency map predicted from algorithm (right) .

or smaller subsets. However, subjects' actual choices were not always consistent with the most non-obvious of their principles.

Mehta, Starmer, and Sugden (1994b) proposed a contrast between "secondary salience" and "Schelling salience". Following Lewis (1969, pp. 24-36), they suggested that when players are not sure what to choose, they choose according to "primary salience", which is "some (possibly stochastic) process that brings one of the labels to the player's mind" (p. 660). Secondary salience is the belief about what creates primary salience for others.

The primary-secondary distinction arises naturally in the SCH model estimated below, although that model was developed to explain behavior in a much wider range of games. In SCH the 'process that brings one of the labels to the player's mind'– its primary salience – is predicted *ex ante* from the SAM model. In the Mehta, Starmer, and Sugden (1994b) paradigms primary salience is measured by having people simple choose objects. Using SAM a primary salience prediction is delivered for all images; no new data or free parameters are needed.

In contrast to primary and secondary salience, an object has Schelling salience if it is unique or is chosen by a rule that leads to unambiguous results. A key distinction is that Schelling salience need not arise from primary or secondary salience. Indeed, Mehta, Starmer, and Sugden (1994b); Mehta, Starmer, and Sugden (1994a) find evidence for both secondary and Schelling salience in their data, where primary salience is calibrated by what people choose in a nonstrategic "pick one" condition.

Focality is likely to work differently in hider-seeker games (HS). In equilibrium,

all agents should choose all locations equally often, regardless of whether they are hiding or seeking. An early study Rubinstein, Tversky, and Heller (1997) (RTH) used a four-choice hider-seeker game. Their canonical example is a choice between four letters ordered from left-to-right, where one letter is a singleton subset, like so:

A B A A

RTH hypothesizes that the left and right A letters are avoided (because of "extremity-aversion"; cf. Bar-Hillel (2015)). They hypothesize that the single B is clearly focal because it is both visually and semantically unique; and it will therefore be avoided by hiders. That leaves the second "interior" A from the right, which is least focal when compared to other choices (and therefore uniquely non-focal, giving it an ironic strategic focality).

In these early studies, extremity-aversion and B-focality are simply intuitions guided by no data. On this basis, RTH predicted that the third A would be chosen most often. Indeed, in their experiments the third A is chosen most frequently both by hiders (40%) and seekers (45%). As a result, there is a "seeker advantage" because the seekers win more often than the prediction by Nash equilibrium. Our replications in Caltech and UCLA subjects found much lower rates of the choice of the third "inner" A, around 29%, closer to the Nash 25% prediction.

Ruma Falk, Raphael Falk, and Ayton (2009) used visual games similar to the four-letter choice. One game required choosing 3 cells out of the 25 locations in a 5x5 matrix. They observe an edge aversion and a seeker advantage.[2]

Our modelling builds upon Vincent P. Crawford and Iriberri (2007a) (hereafter CI), who advanced a novel analysis of games like ABAA, based on level-k modeling. They hypothesized that behavior could be consistent with a level-k approach, [3] in which level-0 behavior is influenced by salience. Specifically, CI assumed that level k type only best responded to level k-1 types, and that the population did not contain any actual level zero types. Under this framework, they estimated both level zero players' preferences towards different options (saliency biases) and population frequencies of level types. The general approach fits behavior well. Our paper

[2]Based on data reported in their paper, the seeking win rate in this experiment is 10.37% while the chance level is only 6.25%, implying a seeker advantage of +4.12%. These numbers are rather close to our own, about 7% and 9%, although the paradigms differ a lot.

[3]See Stahl II and Wilson (1994); Nagel (1995) and see Vincent P. Crawford, Costa-Gomes, and Iriberri (2013) for a thorough review.

expands on this approach by predicting saliency independently of choice, using no new data, in location games.

Hargreaves Heap, Rojo-Arjona, and Sugden (2014) questioned the strength of the CI conclusions on the grounds that the salience of the extreme A's and the unique B were estimated parametrically and not constrained across game structures. They created choice sets with a single "oddity" that is visually or semantically unique (e.g. a list of words which are all diseases except the word "fitness".) They test whether the oddity is equally salient for level 0 players in three types of games– coordination (matching), discoordination (players win if they both choose something different), and hider-seeker. They reject the hypothesis that level 0 salience is the same across games. Vincent P Crawford (2014) comments on their paper. We find better "portability" of salience across matching and hider-seeker games (see section 6).

2.3 The Saliency Attentive Model (SAM) algorithm

Algorithms which take images as inputs, and output predictions about where people look, have been an active area of research since the 1990s. The goal of these algorithms is to predict where neurotypical adult humans will look when they are freely gazing at images.

These algorithms have made great strides. We used a recent algorithm, the Saliency Attentive Model (SAM) to obtain a normalized saliency map for each image (Cornia et al., 2016). The SAM algorithm is based on a machine learning method called a convolutional long short-term memory network (LSTM). The network is trained on the four most popular often-access saliency datasets (SALICON, MIT1003, MIT300, CAT2000) and is trained to best predict actual human gaze patterns recorded from eyetracking. The reported performance of SAM on the website MIT-Saliency is 0.88 using the AUC-Judd measure (Riche et al., 2013), where 1.0 is perfectly accurate. This accuracy significantly surpasses earlier algorithms, and approaches the accuracy of the best human-to-human benchmark[4] (0.92).

The SAM algorithm takes one image as an input and outputs its predicted saliency map. The saliency map is a saliency value from zero to one (least salient to most salient) assigned to each pixel on an image. We adopted the default parameters from

[4]The best human benchmark indicates how much two large different sets of human fixation maps correlate. Each set contains many different individuals. Thus, this limitation is due to individual differences of their fixation patterns (Judd, Durand, and Torralba, 2012).

the original approach and applied it to our image dataset. As a result, the saliency predictions we are using to make predictions of location choices in games have no free parameters. Figure 2.4 is a specific example of the SAM saliency map from one of the pictures we used in our experiments.

A little history of saliency mapping may be useful here to convey how well-founded these algorithms are (and to support further discussion below).

Inspired by a deep understanding of how the human visual system prioritizes attention, a series of progressively improving algorithms were developed to use visual images as inputs, and output predictions about where people will look in the first 1-2 sec of processing (Itti, Koch, and Niebur, 1998; Harel, Koch, and Perona, 2007; Judd, Ehinger, et al., 2009).

These early algorithms used a combination of handcrafted features to extract information about contrast, color, and orientation. Dark-light contrast is special because it marks boundaries between objects. Color and orientation are also thought to have adaptive value in parsing images in ways that are ecologically useful.

These classes of features are called "bottom-up" because they are perceived rapidly, and do not use any abstract information about the meaning or value of what is perceived. "Top-down" perception, in contrast, depends on meaning, personal knowledge and task goals (Frintrop, Rome, and Christensen, 2010).[5]

Consider the stick figure "I". The bottom-up perception is a black vertical line of a certain length, with slightly extended top and bottom horizontal lines on top of the vertical line, surrounded by contrast with a white background. Primates, babies, and adult humans all have the same rapid (< 500msec) perception of this stimulus.

Top-down perception adds meaning, making the perception more behaviorally useful. An English speaker will perceive "I" as a marker of first-person communication; a student just learning Roman numerals will perceive "I" as the number one; and an architecture aficionado may perceive it as an Iconic column, a part of a building. All the latter forms of salience use semantic knowledge – which is local, and acculturated – about the world to inform perception of what "I" means and what to do with that information. Top-down perception of one object rather than another tends

[5]Like many scientific dichotomies, it is hard to draw a sharp line between bottom-up and top-down processes that guide attention (e.g., Awh, Belopolsky, and Theeuwes (2012)). The two processes together can be thought of as a "family of filters" that have been adaptively shaped by forces ranging almost continuously from evolutionarily-conserved universal principles to others locally tuned by personal experience and valuation.

to be influenced what which features are personally relevant, valued, familiar, and novel.[6]

In psychology, attention-getting features that are salient are usefully divided into two categories: Bottom-up and top-down. Bottom-up features are based on sensory physical features such as color, contrast, and orientation. Bottom-up feature salience is devoid of personal meaning and independent of task goals. To the extent that human visual systems work the same way all around the world, the attention-grabbing salience of bottom-up features should be universal. As a result, attentional data from many diverse subject pools can establish what is bottom-up.

The SAM algorithm we use was tuned using human free gaze data on a large number of images, without any special goals or incentives. The subjects are just told to look. These algorithms were **not** designed to predict active choices in games with specific goals, such as matching and hider-seeker. The matching goal, for instance, is to choose a location another person is also likely to choose. This is a top-down influence on perception which is likely to produce visual fixations that are different from free viewing. Thus, the extent to which SAM can predict the influence of predicted salience is probably a lower bound on how well better models, incorporating top-down goals, will do.

2.4 Experimental procedure

Subjects

N=151 subjects participated in total excluding the pilot dataset for power analysis. Among these 151 subjects, N=29 subjects (13 males, 16 females) participated in a lab, one at a time, in a small testing room. The other N=122 subjects participated online through Amazon Mechanical Turk ("mTurk"). This study was pre-registered on the Open Science Framework (https://osf.io/yuqjg/) during data collection and before analysis process. The sample size was pre-determined before the data collection process, based on a pilot study (N= 29) carried out in March 2017.

All N = 29 lab subjects had their gaze data recorded. N=15 subjects were from the Caltech community and N=14 from the neighboring community (there were no differences in results).

[6] See studies on the effects on perception of recent choice history (Awh, Belopolsky, and Theeuwes, 2012), familiarity and novelty (Itti and Baldi, 2009), value for consumer goods (Towal, Mormann, and Koch, 2013), and self-reported "meaning" (Henderson and Hayes, 2017).

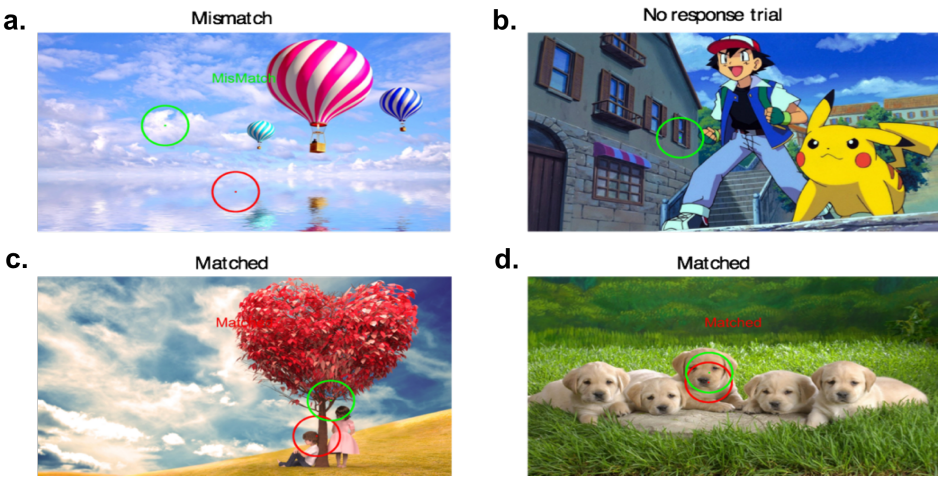

Figure 2.3: Examples of trial outcome.

Experiment designs

The experiment consisted of three blocks of games: matching, the hider-seeker game as a seeker, and the hider-seeker game as a hider. The matching block always came first, followed by the hide and seek blocks in an order randomized between subjects. During each block (game), subjects experienced two sub-blocks. In the "feedback" sub-block the choice the other player made was revealed right after choice. In a "no feedback" sub-block the results were revealed. The matching game consisted of 20 images for each sub-block and the hider-seeker game shared the same set of 19 images for each sub-block. There was unlimited time to read instructions and break but only 6s to make a choice for one trial. The results shown in the feedback condition were drawn from previously tested actual subjects (using different subjects for each image shown). Both feedback and no-feedback blocks were included because each one answers a different interesting question. In ensuring subject comprehension, and especially in testing equilibrium concepts and learning theories, the standard practice is to provide feedback after each trial. However, whether focality is influential even with no feedback about other players' choices is an interesting question too, as in Schelling's thought experiment about the map.

There are some differences between the feedback and no-feedback conditions. The largest is that the matching rate is much higher with feedback (64% vs. 35%). However, the seeker win rate in hider-seeker games is the same in both conditions (9%) and most other differences are not substantial. We therefore report only data from the feedback condition in this main text. The corresponding no-feedback results are in an Appendix.

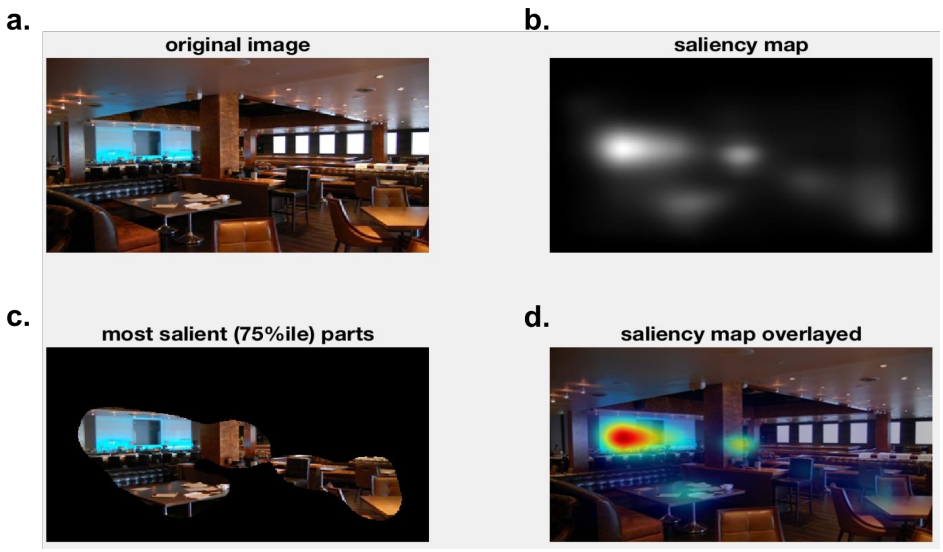

Figure 2.4: A saliency algorithm example.

(a): An original image. (b): The SAM saliency map, in which the brightness indicates salient level. (the lightest point is the most salient spot). (c): The area of original image which is 75% most salient. This area is generated from ranking all saliency values of each pixel. (d): the original image with the saliency heatmap overlaid onto it ("warmer" red colors indicate higher saliency).

Two choices were considered a match if two circles centered at their choice spots (with a radius of 108 pixels ,the circle is about 1/5 of the screen width) overlapped. Figure 2.3 shows examples of result screens that subjects saw during the experiment. The local group and online groups had the same experimental design.

Stimuli

Ninety-two (92) colored visual images were displayed on an eye tracking monitor (1920x1080 resolution). Images were randomly selected from a large image pool (273) with five categories (abstract art, city, faces, society, nature). The image set contained both images with only one obvious saliency spot and more complex images, which have multiple saliency centers (Judd, Ehinger, et al., 2009).

Either is hard for humans to identify any centers. In vision language, the former are called "low entropy" images and the latter are "high entropy" images (examples are shown in a later section). [7]

[7] Entropy is defined by $-\sum p \times log_2 p$, where p is the (0-1) normalized salience of each pixel, and the summation is taken over all pixels on an image. Higher numbers indicate more concentration of salience (entropy is minimized when all values of p are equal). It can be used to characterize the complexity of the image. E.g., a picture with some content on it has less entropy than a white blank paper.

2.5 Analysis and Results

Equilibrium analysis

Equilibrium analysis generates a statistical benchmark for what people might do[8]. The structure of our games deviates from previous matching games, because locations in an x-y space are being chosen. Players click on a pixel, but their choice circle is a 180-pixel circle around their chosen location (an area about the size of an American nickel coin). Matching is defined by circle overlap.[9]

For a matching coordination game, choices by the players of any two different pixels which create overlapping circles constitutes a pure strategy Nash equilibrium. One image contains about two million (=1920×1080) pixels. Any match is a pure equilibrium, so there are an enormous number of equilibria. In addition, there are many mixed equilibria.

For the hider-seeker game, there is a *unique* Nash equilibrium in which all locations are chosen equally often. A special design that, if a circle touches any boundary, it wraps around from the opposite boundary, guarantees the equilibrium.[10]

The proof that equal randomization over all strategies is the unique hider-seeker equilibrium essentially puts logic ahead of bio-logic. The last thing the brain is equipped to do is to *ignore* differences among many objects and choose them equally often. The human perceptual system evolved to efficiently filter a huge amount of information hitting the retina, estimated to be 10^8 bits of information, into only about 100 bits, by focusing attention on only most valuable information —- the most salient locations. For the same reason that we are so good at quickly noticing salient information, we are likely to be naturally bad at rapidly choosing

[8] An important mathematical touchstone is "correlated equilibria" (Aumann, 1974), when both players received public signals and a strategy recommended conditional on different signals. A correlated equilibrium occurs when nobody wants to deviate from recommended strategies. Stop signs and green-yellow-red traffic lights, for example, act as correlating devices to create a commonly-observed visual signal which coordinates traffic and reduces accidents. In these terms, our study is about whether visual salience of image locations works as a correlating device.

[9] The size of the circle around the chosen pixel will influence random matching rates, increasing and decreasing them for larger and smaller circles. This feature could be usefully varied in future research.

[10] For those unfamiliar with game theory, intuition can be gained by a simplified example. Suppose there are just two locations and the hider chooses them with probabilities p and 1-p. If the seeker matches those probabilities, she has a $p^2 + (1-p)^2$ chance of winning. This sum is always lower if the seeker chooses the most likely spot (i.e., the location with $p > 0.5$) because if $p > 0.5$, then $p > p^2 + (1-p)^2$. To defend against this, the hider should mix equally, so $p = 0.5$. Every new location that is added should also have a $\frac{1}{n}$ chance of being chosen (if there are n locations) by an iterated logic.

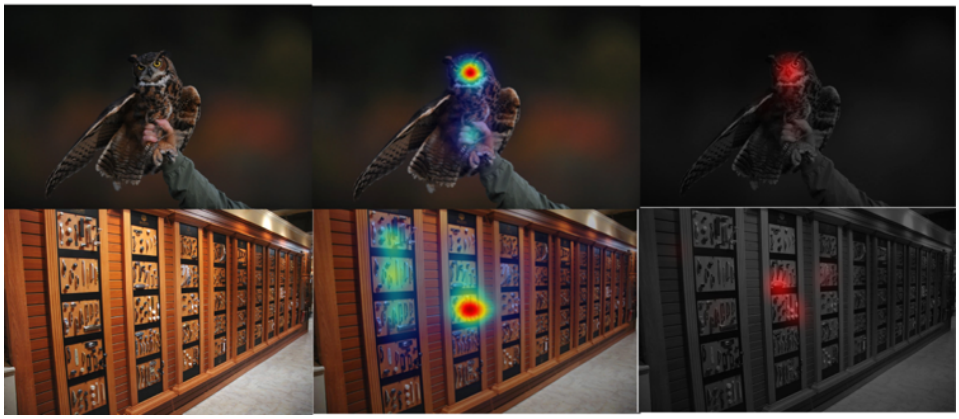

Figure 2.5: Matching game image, saliency heatmap, and choices (red).

(Left column) Two original images. (Middle column) The original images overlaid by the saliency maps.(Right column) The grayscale original image overlaid with the actual empirical choice distributions (more red indicates more frequent choice of an area).

what is unsalient.[11].

Matching game

To analyze the behavioral data, we first test whether subjects are playing an equal random mixture across all pixels (and associated saliency levels). To compare results from different images, all saliency values in this section refer to the normalized levels, which are the rank percentiles of raw measures from the algorithm, ranked within each image. We calculated the normalized saliency value for each click point (the choice saliency level) and then compared these values against the baseline of equal randomization independent of salience. Kolmogorov-Smirnov tests reject the hypothesis of randomness for all treatment conditions ($p < 10^{-4}$).

To get a direct view of how salience affects choices, we plot the choices from all the subjects on two different specific images, shown in figure 2.5.

The saliency heat map (middle column) uses color heatmap to show more salient locations from the SAM algorithm. The right column shows, in redscales, frequencies of subjects' location choices. The predicted saliency in the middle column and

[11] A similar conflict between logic and biology occurs in the games "rock, paper, scissors" and matching pennies (e.g., Vincent P. Crawford, Costa-Gomes, and Iriberri (2013)). When players display the three choices with their hands, there is a slight tendency to match an opponent's choice (e.g., playing rock against rock) more often than predicted in equilibrium. The explanation is that imitation of another person's body movements is such a highly-adapted automatic behavior, that the brain cannot inhibit the response, even though it reduces performance (e.g., you should play paper rather than imitating rock).

the observed choice maps in the right column are highly overlapping.

Statistically, the mean of the saliency levels of actual points chosen in the coordination game is 0.87, which is far beyond chance level 0.5 ($p < 10^{-4}$).

How predictable is the matching rate across images?

An intuitive hypothesis is that the matching rate in an image should be affected by how spread out saliency is in the image. (This prediction also comes structurally from the SCH model described below). When saliency is highly concentrated, then the rate of choosing the same pixels, and thus the matching rate, should increase if people are trying to match. And if there are many salient locations, matching should be lower.

For example, most people should prefer to play the matching game on a white paper with only one star in the center (a single peak of the 2D saliency distribution) instead of a white paper with four identical stars (four peaks of saliency) in four corners.

Figure 2.6 shows that indeed, the matching rate is negatively correlated with the number of saliency centers (Pearson $r = -0.45$, $p = 10^{-4}$).

The actual and predicted matching rates span almost the entire range from perfect matching (100%) to random (7%). These results suggest that for *any* image at all, the matching rate can be predicted with substantial accuracy. Put the other way around, it should not be difficult to find images with saliency distributions that will predictably yield either near-perfect matching or random matching.

Hider-seeker games

For the hider-seeker game, we also present a result of an example image (Figure 2.7) using the same display format as in Figure 2.5. Figures 2.7 show that subjects' choices (c and d) are much more spread out in both hiding and seeking than in the matching game where saliency and choice strongly overlapped.

In Figure 2.7d, the peak of the hider choice distribution no longer falls in the most salient area. However, in this example Figure 2.7d hiding game, the most salient locations are still chosen frequently by hiders (suggesting that, in this example, hiders are not being very strategic).

The general statistics are consistent with the direction of effects in these two examples. In the hiding game, the mean saliency level of user click points is 0.53,

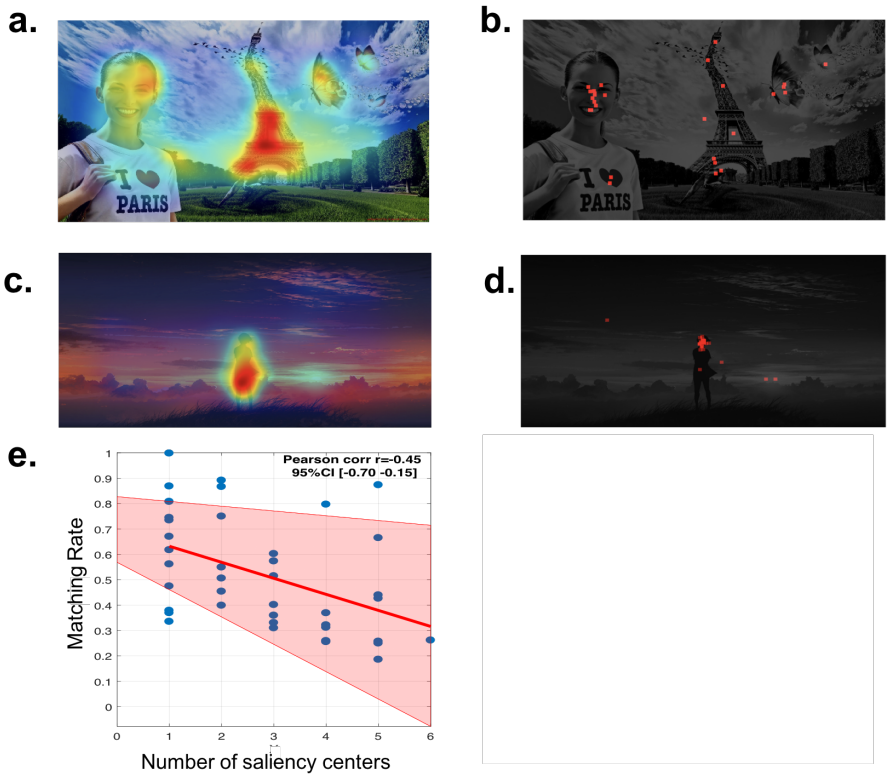

Figure 2.6: Correlation across images between matching rate and number of saliency centers.

(a), (c) are saliency maps of two images. (a) is an example of image with many saliency centers (c) has only one saliency center. (b,d) are corresponding maps (red dots) of actual choice data in a matching game. The choice map in (b) is more dispersed because the saliency centers in (a) are more numerous. (e) is the correlation between the number of saliency centers and the matching rate.

close to chance level.[12] The mean saliency level in the seeking game is 0.61.[13]A replication sample ($N = 29$) with payoffs ten times higher had very similar results, 0.51 and 0.64 for hiders and seekers, respectively.[14] A paired t-test showed that this difference in choice saliency between hiders and seekers is highly significant ($p < 10^{-4}$), reflecting what is seen in figure 2.7. The no-feedback results have a similar difference (see Appendix 2.A).

[12] p-value = 0.02, t-test CI: $[0.51, 0.56]$

[13] mean = 0.61, p-value$< 10^{-4}$

[14] Hiding: p-value for test against null of .50 saliency = 0.59, CI: [0.48,0.54], seeking: p-value $< 10^{-4}$

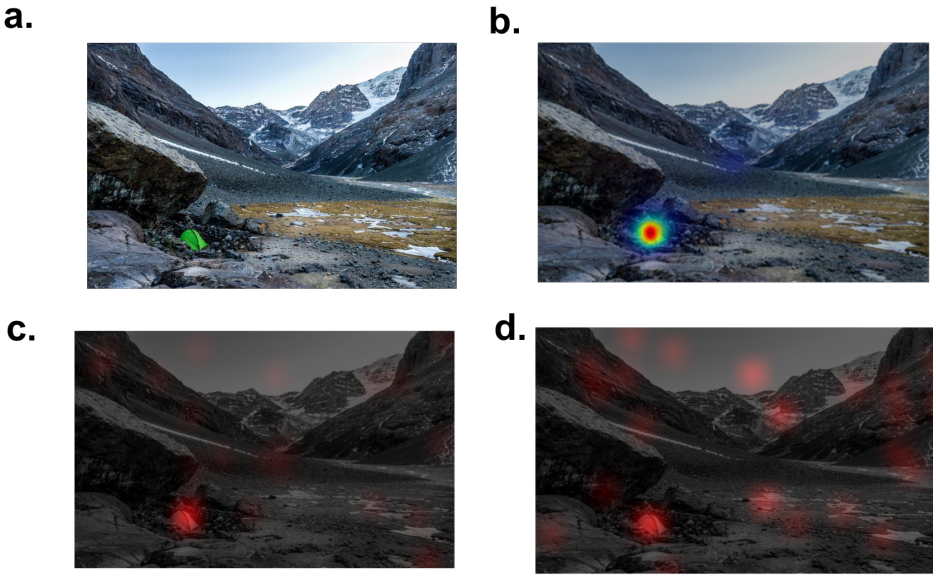

Figure 2.7: Hider-seeker game image, saliency, and choice.

(a): The original image. (b): The original image overlaid by the saliency map. (c,d): The grayscale original image overlaid with the actual empirical choice distributions (more red indicates more frequent choice of an area). c is for seekers' choices and d is for hiders' choices.

Seeker's advantage

In this section, we report the realized matching rate and compare it to the chance outcome. The theoretical value for two circles to match from two random players is four times the area (distance within a larger circle of two times the diameter of the original circle) of the circle used in the experiment over the entire area (1920×1080). That frequency is 0.071.

Table 2.1 presents the realized matching probability in a specific game condition[15].

To check robustness of these results, we also did the hider-seeker game online in a high-payoff condition with payments 10x as large (N=29) and in a between-subjects condition where subjects played only one role (N=53). In both conditions, we got the same seeker win rate of 0.09, exactly as in the main experiment.

The 9% win-rate for seekers does not seem much larger than the equilibrium pre-

[15] Tests to compare the matching rates with random baseline were carried out by bootstrapping a hiding data and a different person's seeking data (or two data points from matching game) for 1000 batches (batch size is total number of different pairs). We get the empirical distribution for the matching rate and statistical significance against baseline 0.071 from that bootstrap. Specifically, each sample is drawn by matching two random users (different ones). The batch seeking win rate is calculated accordingly. All values were calculated from the average of 500 iterations of randomly matching two data points from the data set if two subjects were in the same sub-block, same image.

	Matching rate	N of observations
Nash mixed prediction	0.071	
Matching game	$0.64(4*10^{-4})$	1147
Hider-seeker game	$0.09(3*10^{-4})$	1090
Hider-seeker high payoff	$0.09(6*10^{-4})$	462

Table 2.1: Realized matching rate

Statistical tests against the null hypothesis that the seeker win rate is the baseline level and choices are independently and identically distributed across subjects (which is the Nash benchmark prediction).

diction of 7%. However, in the null hypothesis of Nash equilibrium, this win rate should be identically distributed for all images and for all people. This null hypothesis supplies a lot of statistical power because it justifies pooling all observations together. So when independence is assumed, the test has enough power to establish that 9% is significantly higher than 7%. A more conservative approach averages all data within an image and tests whether the image-wise matching rates are above 7% (N=19, p= 0.0005). A different conservative approach averages win rates for individuals and tests whether the average individual seeker win rate is different than Nash (N = 29, p= 0.002).

2.6 A Saliency-perturbed Cognitive Hierarchy Model (SCH)

This section describes a parametric behavioral model meant to explain choices and their salience-sensitivity, following Crawford and Iriberri (2007a). It uses the level-k model of Stahl and Wilson (1994) and Nagel (1995), later extended by C. F. Camerer, T. Ho, and Chong (2004).

The model combines cognitive hierarchy levels, a quantal response function (softmax), and a saliency-perturbed level 0 assumption.

General Model description

The population consists of different levels of players starting from level zero and with level k players with frequency f(k) (assuming to be Poisson distributed with parameter τ).

For all levels of players, we allow some degree of randomness which will be described using a conventional logit softmax function $\frac{e^{\lambda x_n}}{\sum_m e^{\lambda x_m}}$ with parameter λ. When λ equals zero, agents choose pure randomly without any response to differences in valuation. When λ approaches infinity, agents choose the best option.

In this SCH specification, the nonstrategic level zero players weakly prefer salient choices. Preference is incorporated by higher subjective value assigned to saliency of choices. Formally, the probability of choosing point n depends on the direct saliency value[16] S_n according to

$$P_{0n} = \frac{e^{\lambda(1+\mu S_n)}}{\sum_m e^{\lambda(1+\mu S_m)}}$$

Note that if $\mu = 0$, level 0 types choose randomly among all points. We assume that λ and saliency weight μ are common across subjects, although many heterogeneous versions are conceivable (e.g., Rogers, Palfrey, and C. F. Camerer (2009)).

Other than level zero players, all other levels of players behave in the same way as in CHC. Level k players assume that all other players are only of lower levels (0 to $k-1$), using normalized Poisson frequencies $f(k)$. A level k player calculates the expected payoffs of choosing n, denoted as EU_{kn}. The probability for player i choosing option n is

$$P_{kn} = \frac{e^{\lambda EU_{kn}}}{\sum_m e^{\lambda EU_{km}}}$$

Note that saliency only enters *directly* into the value calculations of level 0 players. The challenge is to see whether saliency which only directly influences the lowest-level 0 players then has an amplified effect on higher-level thinkers, through their beliefs and choices (a la Mehta, Starmer, and Sugden (1994a)), "secondary salience" derived from primary salience).

Model fitting results

Besides the SCH above, there are many other ways to specify models of limited strategic thinking, which have been mixed and matched in previous research. We therefore fit six different model specifications to the hider-seeker data (see Appendix A3).

A simple mixture of Nash play and level 0 fits the least well, but is not bad; it seems that a lot of what is happening in the hider-seeker data is level 0 salience.[17]

[16] We will use the direct saliency measure from the algorithm (raw value) for the modeling input instead of normalized ones (saliency ranks). The difference is that saliency ranks are distributed linearly among points while the raw value is distributed closer to an exponential distribution. In the Q-Q plot, normalized values makes the results comparable with chance diagonal line since rank value is linearly distributed.

[17] Note that level 0 plus Nash is a non-starter for matching games since Nash predicts all pixel choices are equilibria.

-	-	-	Hider	Seeker
Best fit parameters	$\lambda = 100$	$\mu = 0.06$	$\tau_h = 0.4$	$\tau_s = 0.1$
Number of observations			1096	1090
95% CI	[72.3,100]	[0.05,0.08]	[0.32,0.47]	[0.08,0.13]

Table 2.2: Estimation details, role-specific SCH

The parameters μ and λ are constrained to be the same for both hiders and seekers. The confidence interval in the table is calculated using bootstrap method with batch size 1096 for hider, 1090 for seeker and number of iterations 100.

Some specifications restrict the frequency of actual level 0 types to be zero, $f(0) = 0$, as if level 0 players are a figment of the imagination of higher-level types (though see Wright and Leyton-Brown (2019)). Restricting $f(0) = 0$ clearly degrades fit (Table A2). We therefore focus only on $f(0) > 0$.

A close relative of SCH is the "Level k" model, which conventionally refers to level k types believing all others are level k-1 (rather than distributed from 0 to k-1) (Vincent P. Crawford and Iriberri, 2007a; Vincent P. Crawford and Iriberri, 2007b). This is usually estimated non-parametrically, allowing all frequencies $f(k)$ (up to some maximum k) to be estimated separately.

Both SCH and Level k specifications with role-specific level frequencies fit equally well by the AIC criterion (although SCH is a little better by BIC).

We first focus on the best specification of SCH. It has four free parameters– μ, λ, and two role-specific parameters τ_s, and τ_h. (Allowing role-specific λ and μ parameters is fits worse due to the large BIC penalty for extra parameters.)

We used a standard training-testing separation to avoid over-fitting. Recall that each subject does two sessions.[18] We use the first session as a training set to estimate parameters. The parameter values are then fixed and used to predict the choices in the second session test set (see Appendix). The best fitted parameters and measures of fit are shown in Table 2.2.

Figure 2.8 shows the comparison between actual choice density maps versus model predicted density maps for the hider-seeker game. Training data are shown in the top Figures 2.8ab and test data on the bottom Figures 2.8cd. In the choice data plot, there is a sharp density increase starting around 0.9 saliency for both games (although note that the y-axes are different, so the increase is not the same magnitude

[18]Two sessions contain different image sets. A first session of normal payment trials and a second session of high payment trials.

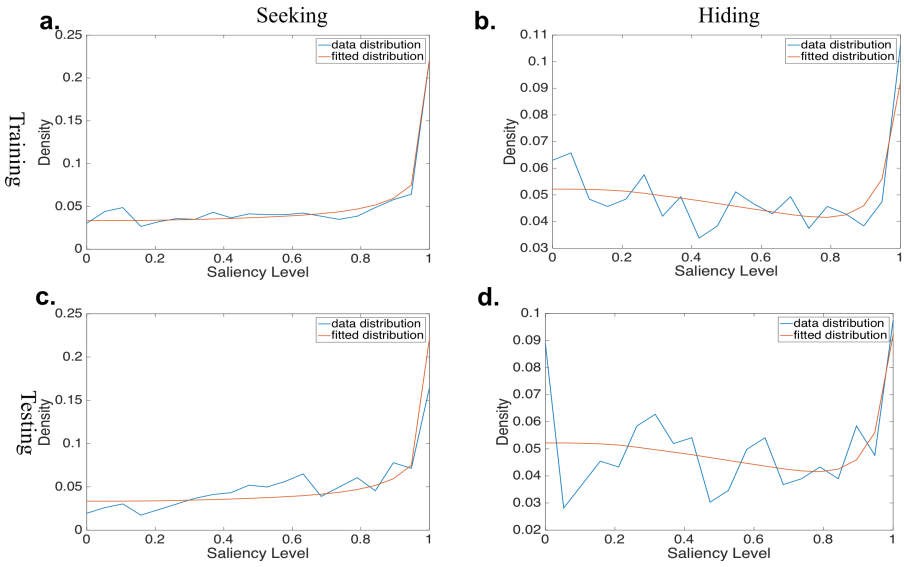

Figure 2.8: Frequency of choice by saliency level with model fitted distributions.

The x-axis is the saliency values of all click points. Each point on a graph indicated what percentage of choices were made for locations within images based on the saliency of those locations. a: choice data and model prediction in the training dataset seeking condition. b: choice data and model prediction in the training dataset hiding condition. c: choice data and model prediction in the testing dataset seeking condition. d: choice data and model prediction in the testing dataset hiding condition.

for both hiders and seekers). There is also a smaller trend of *decreasing* choice from the lowest saliency to 0.9 saliency level for hiders (but not for seekers), reflecting the fact that some hiders choose low-salience locations. SCH can roughly fit these two major features of the data.

The best-fit values of τ, 0.4 and 0.1 for hiders and seekers, are much lower than typical estimates (e.g. Camerer, Ho and Chong, 2004): around τ =1.5. (See also Riche et al. (2013), although Fudenberg and Liang (2018) find minimal prediction error in a large interval (0, 1.25).)

We think the low values of τ estimated in our data for SCH results from the fact that the ability to identify τ is limited in these visual choice games. A much better way to identify τ is with games where different level types choose distinct strategies in a way that is carefully designed to separate them (such as matrix games pioneered by Stahl II and Wilson (1994), also see Nagel (1995); T.-H. Ho, C. Camerer, and Weigelt (1998); Costa-Gomes and Vincent P. Crawford (2006); Kneeland (2015); Fragiadiakis, Kovaliukaite, and Arjona (2017)).

As Figure 2.8 shows, in hider-seeker games the data are close to a random mixture across low saliency levels and a spike at saliency above .90. This can be approxi-

mated rather well by level 0's who exhibit a slight bias toward saliency, and a small percentage of level 1 seekers who approximately best-respond to the highest-saliency location.

The Level-k model can lend some insight here about level frequencies. Compared to SCH, the best Level-k specification estimates lower frequencies of level 0 ($\hat{f_s}(0) = .17$ and $\hat{f_h}(0) = .29$ for seekers and hiders) and a higher salience weight $\hat{mu} = .18$ for those level 0 types. Level-k also estimates larger frequencies of level 2 and 3 types ($\hat{f_{seeker}}(3) = .66$, $\hat{f_{hider}}(2) = .61$). While the overall Level-k fit is just a little less accurate than SCH, this type distribution is more reasonable than the SCH estimates (see Appendix, Level-k Model 6 training–testing comparison for Level-k visual comparison to 2.8 above; the graphs look almost identical). Furthermore, it is clear that a single-peaked SCH with Poisson $f(k)$ cannot approximate the best-fit Level-k distribution frequencies, which are substantial for both level 0 and higher levels. So while it is clear that both specifications fit the saliency-choice profiles adequately (as seen in the Figures), they suggest different evidence of level frequencies. Future research could surely develop better methods, and games in which to more precisely estimate the $f(k)$ frequencies.

Based on the best-fit saliency weight value μ (0.06), the perturbation towards saliency is not that large. It only changes the probability of choosing in the top 5% salient area to 0.14 from the even chance level 0.05 (as blue curves show in Figure 2.9). Though the hiders and the seekers share the same level zero strategy (by assumption), the strategies of higher levels ($k \geq 1$) are affected in different ways: hiders prefer unsalient places while seekers choose the most salient locations more often. Such mixed choice distributions of both level zero and higher-level players are able to explain the trend differences in Figure 2.8 —the seekers' aggregated choice distribution is monotonically increasing in saliency while the hiders' aggregated choice distribution first has a small decrease, then has a sharp increase at the maximal saliency.

Cross-game predictive validation

To further test generalizability of SCH, parameters estimated from fitting the SCH model to hider-seeker data are next used to predict choice behavior in the matching game. There is no guarantee that this cross-game portability will work well. Identification of the saliency weight μ in hider-seeker comes from the level 0's and from a small portion who sometimes try to match the level 0 opponents and sometimes

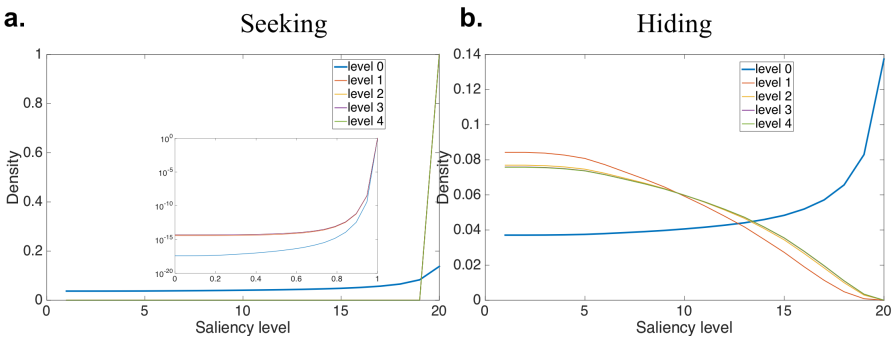

Figure 2.9: Level-specific model predicted choice distribution.

The blue curves on both figures correspond to level zero strategies. In the seeker's case (a), players at level 1 and above highly prefer salient locations. In the larger graph, all levels 1-4 are plotted as green. The inset graph shows there is some variation in sensitivity of level 1 types and above, which can only be seen with a finer-grained y-axis of density. In the hider case (b), level 0 players choose strategies with increase with salience, and all higher-level types choose less salient locations more often.

mismatch. In the matching games, all higher-level types are similarly guided by salience since they are all trying to match the lower-level types. The strength of salience-sensitivity that is estimated in the two cases could easily be different (which is precisely the point of Hargreaves Heap, Rojo-Arjona, and Sugden (2014)).

Figure 2.10 shows the comparison between the experimental data distribution from matching games and model-prediction distributions. The left panel fits hider-seeker data and predict matching data out-of-game. It underpredicts frequency at the highest saliency level. The left panel fits matching data within-game (and will always predict better than between-game). The best-fit $\hat{\tau} = .87$. Deriving data from a different hider-seekr game is not dramatically less accurate than using the same matching game.[19]

2.7 Quantal response equilibrium on saliency

One possible explanation for the saliency bias is that people make mistakes around equilibrium strategies. Such hypothesis can be modeled using a quantal response equilibrium (Palfrey, 2016) with saliency perturbation. Vincent P. Crawford and Iriberri (2007a) has a similar session (section 2) about it and they refer to this model as "equilibrium with payoff perturbations". We will solve for the QRE condition with saliency perturbation in this section and show how QRE can lead to a hider's

[19] We did not do the opposite analysis, predicting hider-seeker from matching, because matching is poorly suited to estimating different levels.

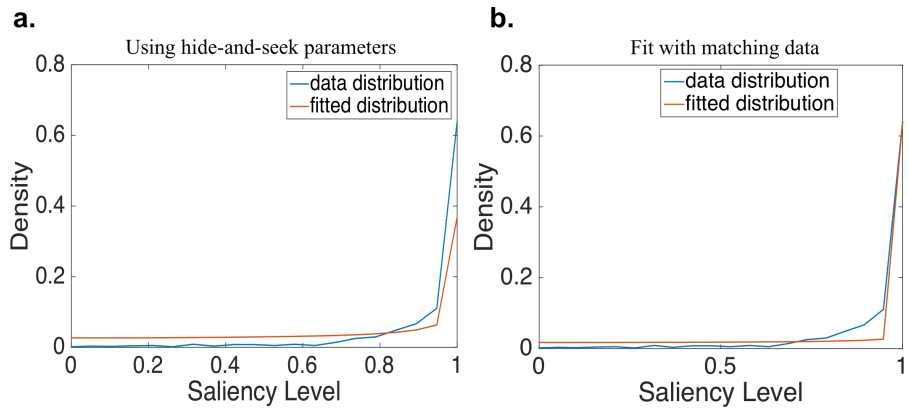

Figure 2.10: The SCH model calibrated on hider-seeker game data can predict matching game choices.

The comparison between the matching data distribution and the fitted matching game distribution. (a) Parameter estimates from the hider-seeker game are used to predict matching game results. log-likelihood: -2176 (b) Parameter estimates from the matching game are used to predict matching game results.. log-likelihood: -1943

advantage.

We make the same assumption as in the SCH model part that each agent has a payoff function for choosing point n in the game:

$$u_n = M + \mu S_n$$

M is the general payoff coming from the game result and μS_n is the saliency perturbation part. Recall that S_n is the saliency level for point(choice) n.

We denote P_{sn} for the probability of choosing n-th point as a seeker and correspondingly, P_{hn} for hider. With the saliency perturbed payoff structure we just mentioned, we can solve for the quantal response equilibrium condition for strategy P^*_{sn} and P^*_{hn} (detailed procedures see Appendix 2.D):

$$\begin{cases} p^*_{sn} = \frac{e^{\lambda(\sigma p^*_{hn} + \mu S_n)}}{\sum e^{\lambda(\sigma p^*_{hm} + \mu S_m)}} \\ p^*_{hn} = \frac{e^{\lambda(1 - p^*_{sn} + \mu S_n)}}{\sum e^{\lambda(1 - p^*_{sm} + \mu S_m)}} \end{cases} \tag{2.1}$$

σ measures the payoff premium: $\sigma = \frac{M_s}{M_h}$.[20]

It will be hard, almost impossible to solve out analytical solutions for P^*_{sn} and P^*_{hn}. And most importantly, it is not that necessary and I will explain why in a minute.

[20]For a hider-seeker game with large search space, usually the seeker needs to have a larger payoff in reality because the chance of successful seeking is small.

Instead of analytically solve the QRE, we can first see how the strategies will change when saliency changes. The result after solving the first order derivative indicates that:

- Seekers always prefer more salient options. Seekers' strategy will be monotonically increasing as saliency level increases for any parameters.
- Hiders' strategies depend on λ. When λ is smaller than a cut-off threshold, hiders also prefer saliency that the probability distribution over all actions monotonically increases as saliency increases. When λ is larger than the cut-off, the probability distribution of the hider's strategy over all actions monotonically decreases as saliency increases.

Figure 2.11 showed the simulation result of QRE strategies for a given set of parameters. The simulation result coincides with what we just analyzed from taking the derivatives that when λ is small, both roles will prefer saliency, and when λ is large, only hider role becomes anti-saliency.

QRE further predicts that seekers' advantage happens only when λ is small, which also means subjects are more randomly choosing than trying to win the game. If they want to win the game and play equilibrium strategies, we should have observed a hider's advantage but we never did.

So is the seeker's advantage that we observed due to small λ? The answer will be no. Recall the realized strategy we plot on figure 2.8 where hider's strategy has a u-shaped property that it first decreases and then sharply increases.

Though the equilibrium solution does not fit our data, it points out a possibility and condition for hider's advantage to happen. This result might be useful in other domains.

2.8 Another two applications of SAM in economics

This section gives out two common applications using SAM in economics. The purpose is to show that SAM bottom-up visual salience is not limited only to the task we specifically designed. Two quite different studies are described: One is the traditional matrix game task; The second is an application to investment experiments, using a combination of actual stock returns and artificial return sequences (Bose et al., 2020).

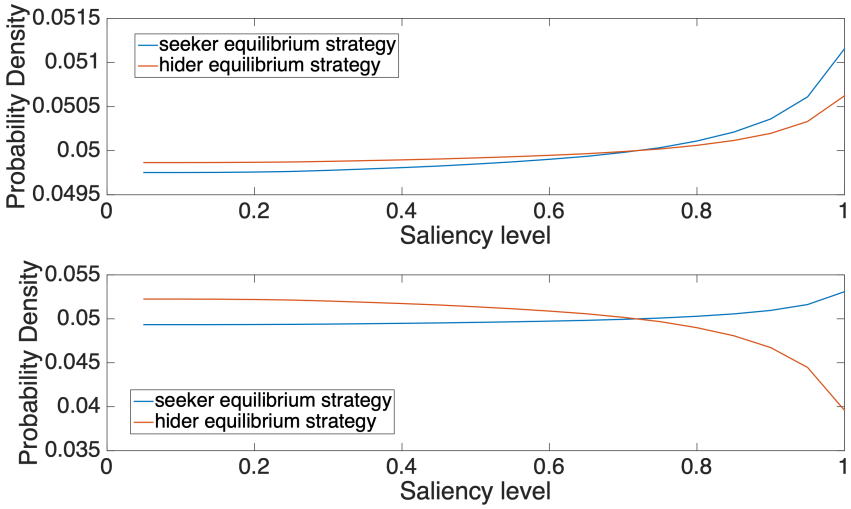

Figure 2.11: QRE prediction of choice saliency map.

The simulation results show how the QRE strategy distributions on saliency domain. The upper figure showed the case when $\lambda < 4.9$ and the lower figure showed the case when $\lambda > 4.9$.

Our application of visual saliency to game theory was motivated by the long-standing interest in focality as a selection principle predicting what happens in coordination games, and using saliency to anchor level-0 behavior to potentially explain the common seeker's advantage in hider-seeker games.

Visual saliency in normal form games

Normal-form games written in matrix format are widely used to study experimental game theory. Several studies have used eye tracking techniques intended to see what we could learn from gaze data in order to further understand the behavior. Little is known about what drives the attention allocation in such typical experimental environments and how different structures of matrix games will affect saliency.

Our visual location game experiments are unusual. Most game theory experiments, following visual conventions in textbook game theory, use normal-form games in a matrix format. To establish boundaries of where visual salience is more useful and it is not, it is useful to ask whether SAM saliency helps explain choices in this widespread visual game format.

Note that the SAM training set does not have any images which even remotely resemble matrices of payoffs. The subjects in matrix game experiments have a clear top-down attentional goal, which is to look at numbers displayed in the visual image in order to make a high-payoff choice. So it would not be surprising if bottom-up

SAM saliency had no predictive power.

To test for possible influence of visual salience, we were lucky to receive data and help from Polonio, Di Guida, and Coricelli (2015). In their experiment, N=56 people played 32 normal-form games with different strategic structures.[21] Eye-tracking was used to record visual attention. We can compare those actual gaze maps with the SAM predictions.[22] Figure 2.12 shows an example and summary statistics. Figure

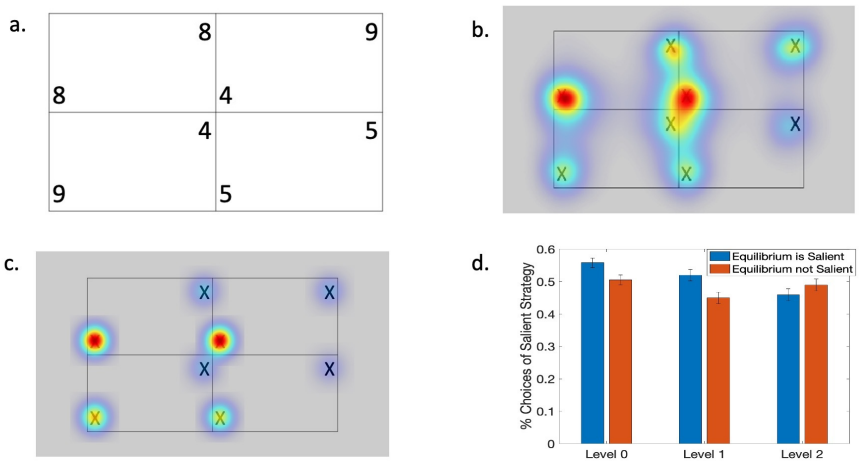

Figure 2.12: Saliency and choices in matrix games.

(a) One (Prison Dilemma) of among all the games used in the experiment. (b) average SAM prediction of all games. (c) Ground truth gaze density map generated by gaze data. (d) Percentage of choices choosing the salient strategy grouped by levels.

2.12a shows *one* example of the kind of matrix subjects see on their computer (it is a prisoners' dilemma in structure). Row player payoffs are in the lower left of each matrix cell, and column player payoffs are in the upper right of each matrix cell.

Figure 2.12b is the average prediction of where people look from the SAM algorithm over all 32 games. There is an obvious top>bottom bias, a bias toward looking at the top row, and a row-player payoff bias. Figure 2.12c is the average measured attention map calculated from gaze data over different types of games (filtering out gazes which are away from payoffs). The comparison between Figures 2.12b-c suggest that the algorithm predicts the actual attention allocation during game plays rather well; this visual impression is supported by conventional statistics used to

[21] Dominant Solvable Self (DSS), Dominant Solvable Other (DSO), Prisoner's Dilemma (PD), and Stag Hunt (SH)

[22] See appendix X for more details.

associate prediction and actual attention in visual science.[23] Much to our surprise, the actual gaze data are quite similar for row players, who choose rows, and column players, who see the same matrix but choose columns. (We quadruple-checked this result.)

The next question is whether bottom-up visual attention, measured by SAM but also highly correlated with recorded eyetracking, influences choices. To test for such an effect on choice, we look at the 24 games which contain a unique equilibrium (an equilibrium row or column). Figure 2.12d presents the different percentage of choices playing equilibrium strategies conditioning on whether the equilibrium strategy is in a more salient location or not. There are three groups of subjects, as grouped by different strategic levels of thinking (based on a sophisticated sequential classification procedure) in the original paper by Polonio, Di Guida, and Coricelli (2015). When the equilibrium choices are in a more salient location, low level players (level 0 and 1) play it more often (p-value = 0.065 for a two-sided test and 0.032 for a one-sided test), while higher level players are not as affected, and play the salience choice a little less often (p-value = 0.57).

2.9 Visual saliency in investment experiments

The last application of visual SAM salience we will describe is in investment experiments (Bose et al., 2020). In their main experiments, participants see many different one-year time series of daily stock prices. For each time series, they decide what proportion of an endowment to invest to earn a return in the next year, after the historical year they have seen. Interested readers can refer to their working paper. Only basic findings are described and illustrated here, for the simpler purpose of understanding whether SAM visual saliency is potentially useful in this different domain of decisions than the location and matrix games described above.

The leftmost figures 2.13 show heatmaps of the density of actual gazes, from N=57 human participants, which were recorded using eyetracking. The corresponding rightmost figures are the time series of actual prices.[24] Blue-to-red heatmap colors show the SAM predictions of where visual attention will be. (Keep in mind that, as

[23] Two validation scores, AUC and CC are commonly used metrics to evaluate how closely saliency algorithm predictions are correlated with actual human gaze in the computer vision field. AUC: area under the receiver operating characteristics curve and CC: Pearson Correlation (see Kummerer, Wallis, and Bethge (2018). The Appendix Table 2.3 shows these statistics.

[24] The eyetracking experiments used a temporally-coarsened series of 60 returns rather than around 250 daily returns, because of the limited spatial resolution of eyetracking. This experiment is just to see how well SAM and eyetracking are correlated, if at all.

	AUC	CC
SAM domain neutral	0.87	0.78
SAM vs fixations(games)	0.96	0.47
SAM vs fixations(price path)	0.81	0.52
Chance level	0.5	0
Range	(0,1)	(0,1)

Table 2.3: AUC measure for SAM on matrix game

The table reports two common evaluation metrics on SAM algorithm, area under the receiver operating characteristics(AUC) and Pearson Correlation Coefficient (CC)(Kummerer, Wallis, and Bethge, 2018). We compare SAM's performance on the domain-neutral images it was trained on against SAM's performance against human eye-fixations for matrix games. The results on both metrics show that SAM predict human fixations far better than chance. Furthermore, AUC reaches a better performance than that tested on neutral images. We think it is due to two reasons: simple image sets and de-noise procedures of the data set.

with all our analyses so far, the SAM predictions are taken directly from implementation of the Cornia et al. (2016) algorithm; there are no extra free parameters.)

In the two examples, salience predicted by SAM appears to be substantially correlated with actual eyetracking gaze density. The pixel-by-pixel Pearson correlation averaged across all time series has an average of .52 (compared to the original SAM-groundtruth correlation which is .78), see summaries in Table 2.3. Participants are often looking where the SAM algorithm predicts they will look.

But looking is not choosing. One can intuitively think of visual perception as necessary, but not sufficient, as a basis for choice. To see whether salience seems to influence choice, the authors compute weights for price-to-price percentage returns based on salience of adjacent prices. These saliency weights are then used to compute a saliency-weighted expected return. For example, in Figure 2.13(b), the SAM salience is highest in the center of the price time series, where there is a big positive return. If salience is guiding value perception, this stock will have a high saliency-weighted expected return and– if influenced by salience– participants will invest heavily in it.

Indeed, there is a substantial correlation between saliency-weighted returns and investment in three experiments. These correlations are above and beyond control variables which include the typical moments of historical returns (mean, standard deviation, and skewness) that also strongly influence investment decisions.

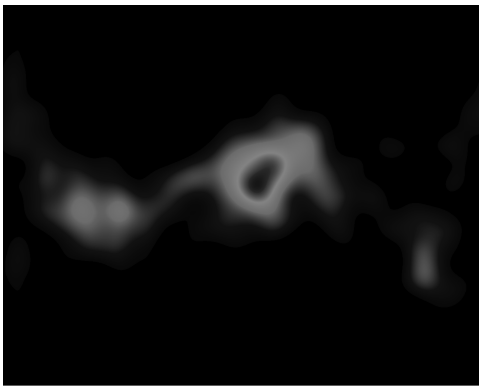

(a) Eye-Tracking Heatmap.

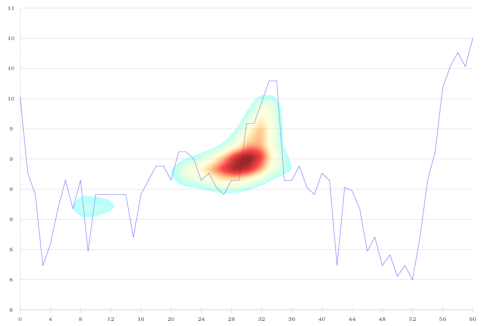

(b) SAM prediction

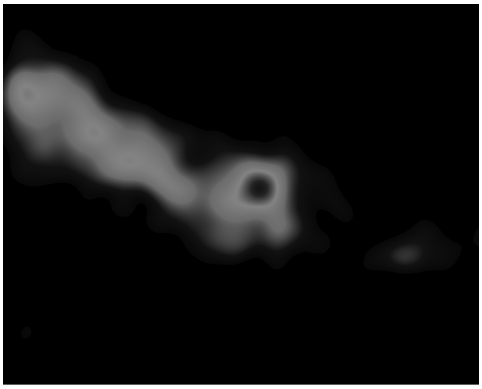

(c) Eye-Tracking Heatmap

(d) SAM prediction

The figure shows two examples of the groundtruth density maps from our experiment on the left, and the SAM-predicted heatmaps where we also show the underlying price paths participants looked at on the right.

Figure 2.13: Heatmaps of price path.

It is sensible to ask whether visual salience corresponds to particular micro-features of the stock price time series, such as large "jumps", short periods of rapid up-and-down negative autocorrelation (generating high visual contrast), primacy or recency of early or later returns, etc. The authors find that about 20% of the variance in saliency values can be explained by such statistics. Saliency – trained on free gaze of a large corpus of images (*none* of which are stock price charts – is responding to some of these statistics, but most of saliency is explained by something else.

2.10 Conclusion and discussion

In our two-player games, players choose locations in images to coordinate (match) or compete (hider-seeker). An algorithm from visual neuroscience (SAM) was used to predict saliency independent of choices. SAM specifies salience precisely with no free parameters for any image. Its generality, and accuracy, means we can skip the common design step of starting simple by using displays like object sets in which saliency is carefully controlled. The method can be immediately applied to more complex images such as Schelling's map and Kenyan bracelets signaling virtuous deworming. It is a use case of how neuroscience of vision might help answer questions of long-standing interest in an important area of economics, namely game theory.

Applying the SAM algorithm solves three problems. First, it predicts with substantial accuracy how often people actually match on different images. Second, when predicted salience is embedded in a cognitive hierarchy or level-k model, the model generates a seeker's advantage. Salience strength estimated from hider-seeker games can also "portably" predict choices in matching games which have a different structure. Third, using eyetracking data and salience, an HMM predicts numerically what is likely to happen under time pressure. The predicted seeker win rate under 2s pressure is 14%, which is not too far from the observed win rate with feedback (15%), and is higher than the 6s win rate of 9%.

Salience in psychology and economics

Our analysis used visual salience in location-choice games. This type of rapid bottom-up visual salience is well-understood. But salience, particularly as it is likely to be important in economics, is more general than that. This brief discussion introduces important properties of salience, as understood in cognitive neuroscience, that could be useful for thinking about how salience is used in economics.

We first note that salience in economics sometimes refers to immediate, bottom-up perception which is likely to be predicted accurately by an algorithm like SAM. For example, there are increased grocery sales of items that are placed at eye level in a store (Dreze, Hoch, and Purk, 1994). Positions on lists and ballot order (in voting) favor items at the top (e.g. Abaluck and Adams (2017); D. E. Ho and Imai (2006)).

Digging deeper, neuroscientists seek to understand a biological activity, such as salience, at three levels: What is the activity for, functionally (Why?); How does it work algorithmically (What?); and How is it implemented biologically in flesh and blood (How?).

What salience is for is clear: Its functional adaptive purpose is to sift through an enormous amount of incoming environmental information to pick out what is most useful. In rapid visual attention, the human retina takes in 10^8 bits of information, ignores almost all of it, and retains just the barest gist such as outlines of objects (100 bits).

The fact that the function of salience is prioritization of valuable information should be appealing to economists. It is also how many modern attention scientists think.

For example, Gottlieb (2012) distinguishes salience for the three different purposes of action, learning, and liking. To illustrate, suppose stimulus features are numerically encoded by their low and high potential reward value. Then salience for action, for example, will direct attention quickly to the highest and lowest value features, because those are the features that are most valuable for a person's action– knowing where to seek reward and where to avoid low reward.

Salience for action roughly matches the hypothesis of "extremity salience" of Bordalo, Gennaioli, and Shleifer (2012). They hypothesized that salience is extra attention given to the most extreme values in a feature comparison, such as the highest and lowest prices relative to a reference price (with further restrictions on salience functions, which are psychophysically well-established, to produce sharper predictions).

Many of the features BGS propose to be extremity-salient are features that naturally map onto the high-low reward variable that Gottlieb's framework rests on. Extremity-salient features include asset payoffs, attributes of goods such as price and quality, and high and low risky outcomes (cf. Bordalo, Gennaioli, and Shleifer (2012); Bordalo, Gennaioli, and Shleifer (2013b); Bordalo, Gennaioli, and Shleifer (2013a)). The BGS formulation therefore fits with Gottlieb's proposal about

how salience guides action (though salience has other purposes).[25]

The BGS specification encodes salience for decision. The two other function of salience are for learning, which prioritizes uncertainly-rewarding novel features, and for liking, which prioritizes known valued features. These could be also included in portable models like BGS is when there is a phenomenon that might be explained by doing so (such as demand for new consumer products).

A different, and prominent, use of salience in economics is that price and value components that are presented to sensory systems, such as explicit price tags the eye can see, receive more decision weight than components that need to be imagined and computed. Beginning in marketing 40 years ago (Russo 1977), many studies have found that unit costs, taxes, shipping costs, and car tolls are given more weight in decision making when they are explicitly presented (Ott and Andrus, 2000; Hossain and Morgan, 2006; Min Kim and Kachersky, 2006; Finkelstein, 2009; Taubinsky and Rees-Jones, 2017). No fancy psychology is needed here. Recognition is easier than recall. Sensory attention is 'better' than memory: it is mentally cheaper, faster, and more accurate.

What's next?

Attention researchers in cognitive neuroscience know a lot about how attention and salience work (particularly at short time scales). Economists do not know as much about mechanistic details, but have become interested in salience and attention. Their usage of the term salience is almost always referring to hidden (and unmeasured) variation in likely salience inferred from observable choices (e.g. Gabaix (2017)).

This reduced-form approach is a big step and has already proved fruitful. However, a better understanding of attentional mechanisms is likely to be useful to understand differences in inferred attention. For example, a clever experimental tax variation study (Taubinsky and Rees-Jones, 2017) estimated that inferred attention to undisplayed taxes was only 25% of the actual tax rate. Chetty, Looney, and Kroft (2009) estimated the same parameter as 35% for groceries in an experiment and 6% for alcohol in field data. These estimates and others in Gabaix (2018, Table 1) vary a lot.

Some variation in these revealed-attention estimates is due to differences in esti-

[25] Kőszegi and Szeidl (2012) hypothesize a similar effect, that choice attributes with higher utility variation will get larger "focus" weight. We do not know whether there is a counterpart to this hypothesis in attention science.

mation methods, data, statistical power, and so on. But some of that variation also could be due to differences in how attention is actually deployed. More knowledge of how attention actually works could then help us understand variation in estimates and make better cross-domain predictions. Better attentional knowledge and models will also predict new causal effects, such as how much exogeneous capture of attention influences choices and what effect time pressure has (as shown by our HMM results).

References

Abaluck, Jason and Abi Adams (2017). *What do consumers consider before they choose? Identification from asymmetric demand responses*. Tech. rep. National Bureau of Economic Research.

Aumann, Robert J. (1974). "Subjectivity and correlation in randomized strategies". In: *Journal of Mathematical Economics* 1(1), pp. 67–96.

Awh, Edward, Artem V. Belopolsky, and Jan Theeuwes (Aug. 2012). "Top-down versus bottom-up attentional control: a failed theoretical dichotomy". In: *Trends in Cognitive Sciences* 16(8), pp. 437–443. ISSN: 1364-6613. DOI: 10.1016/j.tics.2012.06.010. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3426354/ (visited on 05/30/2018).

Bacharach, Michael (1993). "13 variable universe games". In: *Frontiers of Game Theory*, p. 255.

Bacharach, Michael and Michele Bernasconi (1997). "The variable frame theory of focal points: An experimental study". In: *Games and Economic Behavior* 19(1), pp. 1–45.

Bar-Hillel, Maya (2015). "Position effects in choice from simultaneous displays: A conundrum solved". In: *Perspectives on Psychological Science* 10(4), pp. 419–433.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2012). "Salience theory of choice under risk". In: *The Quarterly Journal of Economics* 127(3), pp. 1243–1285.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2013a). "Salience and asset prices". In: *American Economic Review* 103(3), pp. 623–28.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2013b). "Salience and consumer choice". In: *Journal of Political Economy* 121(5), pp. 803–843.

Bose, Devdeepta et al. (2020). "Decision weights for experimental asset prices based on visual salience". In: *Available at SSRN 3654021*.

Camerer, Colin F., TeckHua Ho, and Juin–Kuan Chong (2004). "A cognitive hierarchy model of games". In: *The Quarterly Journal of Economics* 119(3), pp. 861–898.

Chetty, Raj, Adam Looney, and Kory Kroft (2009). "Salience and taxation: Theory and evidence". In: *American Economic Review* 99(4), pp. 1145–77.

Cornia, Marcella et al. (2016). "Predicting human eye fixations via an LSTM-based saliency attentive model". In: *arXiv preprint arXiv:1611.09571*.

Costa-Gomes, Miguel A. and Vincent P. Crawford (2006). "Cognition and behavior in two-person guessing games: An experimental study". In: *American Economic Review* 96(5), pp. 1737–1768.

Crawford, Vincent P (2014). "A comment on how portable is level-0 behavior? A test of level-k theory in games with non-neutral frames by Heap, Rojo-Arjona, and Sugden". In:

Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri (2013). "Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications". In: *Journal of Economic Literature* 51(1), pp. 5–62.

Crawford, Vincent P. and Nagore Iriberri (2007a). "Fatal attraction: Salience, naivete, and sophistication in experimental" hide-and-seek" games". In: *American Economic Review* 97(5), pp. 1731–1750.

Crawford, Vincent P. and Nagore Iriberri (2007b). "Level–k Auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private–value auctions?" In: *Econometrica* 75(6), pp. 1721–1770.

Dreze, Xavier, Stephen J. Hoch, and Mary E. Purk (1994). "Shelf management and space elasticity". In: *Journal of Retailing* 70(4), pp. 301–326.

Falk, Ruma, Raphael Falk, and Peter Ayton (2009). "Subjective patterns of randomness and choice: Some consequences of collective responses." In: *Journal of Experimental Psychology: Human Perception and Performance* 35(1), p. 203.

Finkelstein, Amy (2009). "E-ztax: Tax salience and tax rates". In: *The Quarterly Journal of Economics* 124(3), pp. 969–1010.

Fragiadiakis, D., Ada Kovaliukaite, and D. Rojo Arjona (2017). *Testing cognitive hierarchy assumptions*. Tech. rep. Working paper, 2017. 4, 13, 14, 15, 17.

Frintrop, Simone, Erich Rome, and Henrik I. Christensen (2010). "Computational visual attention systems and their cognitive foundations: A survey". In: *ACM Transactions on Applied Perception (TAP)* 7(1), p. 6.

Fudenberg, Drew and Annie Liang (2018). "Predicting and Understanding Initial Play". In:

Gabaix, Xavier (2017). *Behavioral inattention*. Tech. rep. National Bureau of Economic Research.

Gottlieb, Jacqueline (2012). “Attention, learning, and the value of information”. In: *Neuron* 76(2), pp. 281–295.

Harel, Jonathan, Christof Koch, and Pietro Perona (2007). “Graph-based visual saliency”. In: *Advances in Neural Information Processing Systems*, pp. 545–552.

Hargreaves Heap, Shaun, David Rojo-Arjona, and Robert Sugden (2014). “How Portable Is Level–0 Behavior? A Test of Level–k Theory in Games With Non–Neutral Frames”. In: *Econometrica* 82(3), pp. 1133–1151.

Henderson, John M. and Taylor R. Hayes (2017). “Meaning-based guidance of attention in scenes as revealed by meaning maps”. In: *Nature Human Behaviour* 1(10), p. 743.

Ho, Daniel E. and Kosuke Imai (2006). “Randomization inference with natural experiments: An analysis of ballot effects in the 2003 California recall election”. In: *Journal of the American Statistical Association* 101(475), pp. 888–900.

Ho, Teck-Hua, Colin Camerer, and Keith Weigelt (1998). “Iterated dominance and iterated best response in experimental" p-beauty contests"”. In: *The American Economic Review* 88(4), pp. 947–969.

Hossain, Tanjim and John Morgan (2006). “... plus shipping and handling: Revenue (non) equivalence in field experiments on ebay”. In: *Advances in Economic Analysis & Policy* 5(2).

Itti, Laurent and Pierre Baldi (2009). “Bayesian surprise attracts human attention”. In: *Vision Research* 49(10), pp. 1295–1306.

Itti, Laurent, Christof Koch, and Ernst Niebur (1998). “A model of saliency-based visual attention for rapid scene analysis”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(11), pp. 1254–1259.

Judd, Tilke, Frédo Durand, and Antonio Torralba (2012). “A benchmark of computational models of saliency to predict human fixations”. In:

Judd, Tilke, Krista Ehinger, et al. (2009). “Learning to predict where humans look”. In: *Computer Vision, 2009 IEEE 12th international conference on*. IEEE, pp. 2106–2113. ISBN: 1-4244-4420-9.

Kneeland, Terri (2015). “Identifying higher–order rationality”. In: *Econometrica* 83(5), pp. 2065–2079.

Kőszegi, Botond and Adam Szeidl (2012). “A model of focusing in economic choice”. In: *The Quarterly Journal of Economics* 128(1), pp. 53–104.

Kummerer, Matthias, Thomas SA Wallis, and Matthias Bethge (2018). “Saliency benchmarking made easy: Separating models, maps and metrics”. In: *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 770–787.

Lewis, David (2008). *Convention: A philosophical study*. John Wiley & Sons.

Mehta, Judith, Chris Starmer, and Robert Sugden (1994a). "Focal points in pure coordination games: An experimental investigation". In: *Theory and Decision* 36(2), pp. 163–185.

Mehta, Judith, Chris Starmer, and Robert Sugden (1994b). "The nature of salience: An experimental investigation of pure coordination games". In: *The American Economic Review* 84(3), pp. 658–673.

Min Kim, Hyeong and Luke Kachersky (2006). "Dimensions of price salience: a conceptual framework for perceptions of multi–dimensional prices". In: *Journal of Product & Brand Management* 15(2), pp. 139–147.

Nagel, Rosemarie (1995). "Unraveling in guessing games: An experimental study". In: *The American Economic Review* 85(5), pp. 1313–1326.

Ott, Richard L. and David M. Andrus (2000). "The effect of personal property taxes on consumer vehicle-purchasing decisions: a partitioned price/mental accounting theory analysis". In: *Public Finance Review* 28(2), pp. 134–152.

Palfrey, Thomas R. (2016). *Quantal Response Equilibrium-a Stochastic Theory of Games*. Princeton University Press. ISBN: 0-691-12423-X.

Polonio, Luca, Sibilla Di Guida, and Giorgio Coricelli (2015). "Strategic sophistication and attention in games: An eye-tracking study". In: *Games and Economic Behavior* 94, pp. 80–96.

Riche, Nicolas et al. (2013). "Saliency and human fixations: state-of-the-art and study of comparison metrics". In: *Computer Vision (ICCV), 2013 IEEE International Conference on*. IEEE, pp. 1153–1160. ISBN: 1-4799-2840-2.

Rogers, Brian W., Thomas R. Palfrey, and Colin F. Camerer (2009). "Heterogeneous quantal response equilibrium and cognitive hierarchies". In: *Journal of Economic Theory* 144(4), pp. 1440–1467.

Rubinstein, Ariel (1991). "Comments on the interpretation of game theory". In: *Econometrica: Journal of the Econometric Society*, pp. 909–924.

Rubinstein, Ariel, Amos Tversky, and Dana Heller (1997). "Naive strategies in competitive games". In: *Understanding Strategic Interaction*. Springer, pp. 394–402.

Schelling, Thomas C. (1980). *The strategy of conflict*. Harvard University Press. ISBN: 0-674-84031-3.

Stahl II, Dale O. and Paul W. Wilson (1994). "Experimental evidence on players' models of other players". In: *Journal of Economic Behavior & Organization* 25(3), pp. 309–327.

Taubinsky, Dmitry and Alex Rees-Jones (2017). "Attention variation and welfare: theory and evidence from a tax salience experiment". In: *The Review of Economic Studies* 85(4), pp. 2462–2496.

Towal, R. Blythe, Milica Mormann, and Christof Koch (2013). “Simultaneous modeling of visual saliency and value computation improves predictions of economic choice”. In: *Proceedings of the National Academy of Sciences* 110(40), E3858–E3867.

Wright, James R and Kevin Leyton-Brown (2019). “Level-0 models for predicting human behavior in games”. In: *Journal of Artificial Intelligence Research* 64, pp. 357–383.

APPENDIX

2.A Results from no-feedback trials

The realized matching rate when there is no feedback result shown throughout the experiment is as follows:

Figure 2.A.2 below is a quantile to quantile plot, plotting the percentage rank of saliency for each location against the percentage rank of choice frequencies for those locations. To get the Q-Q plot, we first mapped all users' choice data (not only click points, but all points which fell into the circle) onto a one-dimensional saliency value, normalized from zero to one. (The highest saliency point in each entire image is one, and the lowest is zero). Then we ranked all these realized saliency values for all choices in the targeted sub-block. We also transformed the rank of the choice frequencies across all subjects into rank percentages. We plotted the normalized saliency value, which was also the percentage of saliency, against the percentage of points chosen with the same saliency ranking. The Q-Q plot below shows that all quantiles of choice data are above the same quantiles of saliency level, and hence above the diagonal dashed line that would result if people were choosing independently of saliency.

Figure2.A.1 presents both Q-Q plots and density maps in the hiders-seeker game. Figures 2.A.1 a-b indicate that seekers' choices are more biased towards salient locations than hiders' choices are, and both are much less saliency-biased than in the matching games (recall Figure 2.A.2). Keep in mind, however, that the hiders should be choosing locations as low in salience as they can perceive (i.e., a best-response Q-Q curve would be underneath the 45-degree identity line).

The density maps in Figures 2.A.1 c-d take every location in every game, and assign

	No feedback	Number of observations
Nash mixed prediction	0.071	
Matching game	$0.35(4*10^{-4})$	1147
Hider-seeker game	$0.09(3*10^{-4})$	1090
Hider-seeker high payoff	–	462

Table 2.A.1: Realized matching rate for no feedback group.

Statistical tests against the null hypothesis that the seeker win rate is the baseline level and choices are independently and identically distributed across subjects (which is the Nash benchmark prediction).

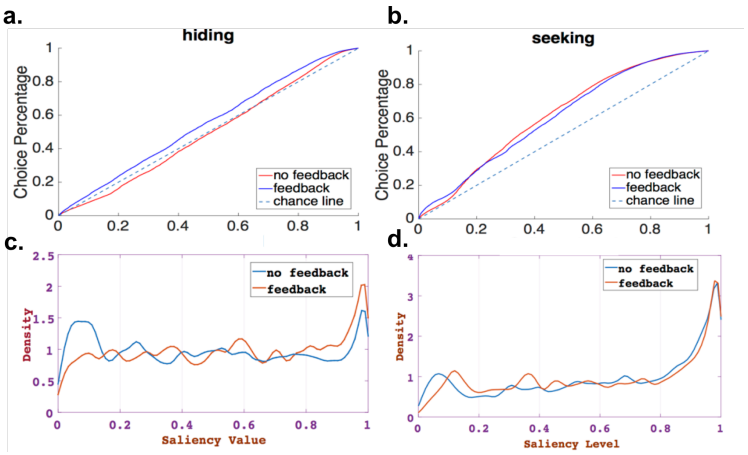

Figure 2.A.1: Hider-seeker game Q-Q plot of choice frequency (x-axis) and saliency ranks (y-axis)

(a, b): Q-Q plots for hiding role (a) and seeking role (b). (c, d): kernel pdf density map of the choice frequency as a function of location saliency ranks. The x- axis is the rank of the saliency values and the y-axis is the probability density. Note: The kernel is Gaussian. The bandwidth is calculated using the formula: $\sigma \times \frac{4}{3N}^{0.2}$, in which is the standard deviation of the samples and N is the number of observations.

each one a saliency level (0-1 normalized within each image), and compute the frequencies with "strategies" (=locations) chosen across all games and subjects. For hider-seeker games, these should be flat horizontal lines in equilibrium (except for sampling error). However, there are a disproportionate number of choices of high-saliency locations (that is, the densities turn up sharply at the right end of the scale). Seekers choose the highest-saliency locations about three times as often, and hiders choose them about two times as often. There is a slightly disproportionate tendency to choose the lowest saliency locations (near zero at the left end of the scale), especially for hiders.

2.B Model comparison

In this subsection, we are going to compare four different sub-models. We choose the Bayesian information criterion (BIC) to be the criterion for model selecting, since it balances between the goodness of fit and the possibilities of overfitting.

In all cases we restricted the softmax sensitivity parameter λ from 0 to 100. Larger values carry little information since $\lambda = 100$ is close to best response. Constraining $lambda$ also makes it easier to create a bootstrapped confidence interval, which is useful due to the non-smoothness of the target function.

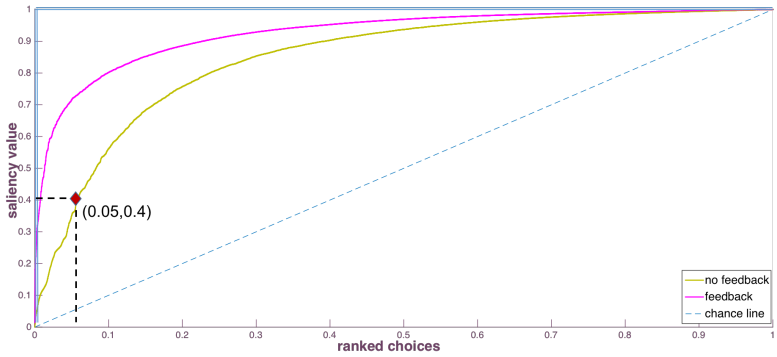

Figure 2.A.2: Matching game Q-Q plot of choice frequency (x-axis) and saliency ranks (y-axis)

The red-diamond point (0.05,0.4) indicates that only 15 percent of choice points were made at the locations at or below 40% salience. Equivalently, 85% of the points fall within the top 60% most salient points. Choices generated by chance would thus correspond to a diagonal line of this plot from $(0,0)$ to $(1,1)$. The maximal accuracy is the blue line: $y = 1$ for all $x > 0$, which means that all choices fall on exactly the most salient point.

Here are descriptions of models we are going to test:

- Model 1: There are only two types of players: 1) naive players who play as level zero players described in the main text. 2) equilibrium players who do pure randomization. Both the proportion of naive players, p_s and p_h serve as free parameters.
- Model 2: There is no level zero player in real population, but in higher level people's belief, there is. The hiders group and the seekers group have different τs but the same μ and λ.
- Model 3: Same setting as model 2, except that level zero players exist both in the belief structure and in the population.
- Model 4: Fit hiding data and seeking data separately using two sets of parameters. Each game has three parameters: μ, λ, and τ. The best fit model of it dominates the best fit model 3 since model 3 is a special case of model 4. However, model 4 allows more free parameters, which BIC value will penalize for.
- Model 5: Fit hiding data and seeking data using a common μ,λ but use level-k framework [26] rather than CH, assuming population is consists of players

[26] see Nagel (1995); Crawford and Iriberri (2007a); Crawford and Iriberri (2007b)

Model	Description	Free parameters (Estimated)	AIC	BIC
1	Level 0+equilibrium	p_s,p_h[a] [1, .3]	12716	12728
2	Role-specific τ_x, f(0)=0	μ,λ,τ_s,τ_h [.004, 99, .46, .002]	12780	12803
3	Role-specific τ_x, f(0)≠0	**μ,λ,τ_s,τ_h [.06, 100, .40, .10]**	**12650**	**12673**
4	Role-specific τ_x, μ_x, λ_x	μ_s,λ_s,τ_s,μ_h,λ_h,τ_h [.01, 90, .40, .07, 90, .50]	12646	12680
5	Level-k role-specific f(k), f(0)=0	μ,λ,$f_s(1)$,$f_s(2)$,$f_s(3)$,$f_h(1)$, $f_h(2)$,$f_h(3)$ [1, 99, .22, 0, .78, .83, .05, .12]	12681	12738
6	Level-k role-specific f(k), f(0)≠0	**μ,λ,$f_s(0)$,$f_s(1)$,$f_s(2)$,$f_s(3)$, $f_h(0)$,$f_h(1)$, $f_h(2)$,$f_h(3)$ [.18, 99, .29 .05, 0, .66, .17, .22, .61, 0]**	**12652**	**12709**

Table 2.B.1: Model comparisons for hider-seeker game

Each model in the table is specified in the list before. BIC is defined as -2*logL + numParam*log(numObs) and AIC is -2*logL + 2*numParam

[a] Stands for proportion of level 0 players

whose level ranging from one to four.

- Model 6: Same setting as model 5, except that model 6 allow level zero type exists.

Table 2.B.1 lists the best fit results of each model. Both BIC and AIC indicate that model 3 is the best performing model. Model 2 performs least for the reason that without level zero specified, such model structure will predict pure anti-salient hiders, which is the opposite of what we observed in data.

It is also not the case that level zero behavior is the only source that drives the fit. The fact that model 1 cannot explain the anti-saliency choices by the hiders serves as one explanation to that. Model 4 performs only slightly worse than model 3 and the best-fit parameters are close to those in model 3, while it contains more parameters. Model 5 also provides a good-fit (same log-likelihood as model 3), but it requires more parameters, which is penalized by the BIC criterion.

The Level-k Model 6 is almost as accurate by AIC and BIC, and we commented on what can be learned from it in the text. Figure 2.B.1 plots predictions of that model

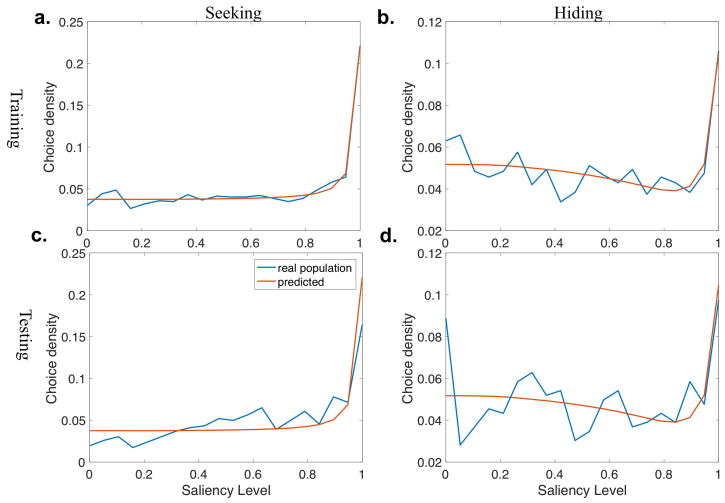

Figure 2.B.1: Level-k Model 6 training-testing comparison

The x-axis is the saliency values of all click points. Each point on a graph indicated what percentage of choices were made for locations within images based on the saliency of those locations. (a): choice data and model prediction in the training dataset seeking condition. (b): choice data and model prediction in the training dataset hiding condition. (c): choice data and model prediction in the testing dataset seeking condition. (d): choice data and model prediction in the testing dataset hiding condition.

and the data, for comparison to text 2.8.

2.C Details of SCH model fitting and SCH converging properties

Method for the model estimation

We used log-likelihood as our target loss function to minimize. Unlike in traditional maximum likelihood estimation, the loss function for our model does not have a closed-form solution for the global minimum. We therefore used a numerical gradient descend algorithm to search for a set of best-fit parameters. In order to avoid finding local optima, the model was initialized from multiple locations and we chose the best one from those multiple initializations. In addition, we constrained the softmax parameter λ to be bounded above by 100. Parameter values above this bound carry less information because they are so close to best response, and bounding λ facilitates bootstrapping confidence intervals.

We use both hiding data and seeking data from the normal pay sessions, maximizing the summation of two log-likelihoods. To avoid overfitting, we tested the model on the second session in which people get higher payment than the first session. It is not the ideal training/testing separation since the payment is different. But due to the fact that cognitive hierarchy model could be sensitive to interdependence at the

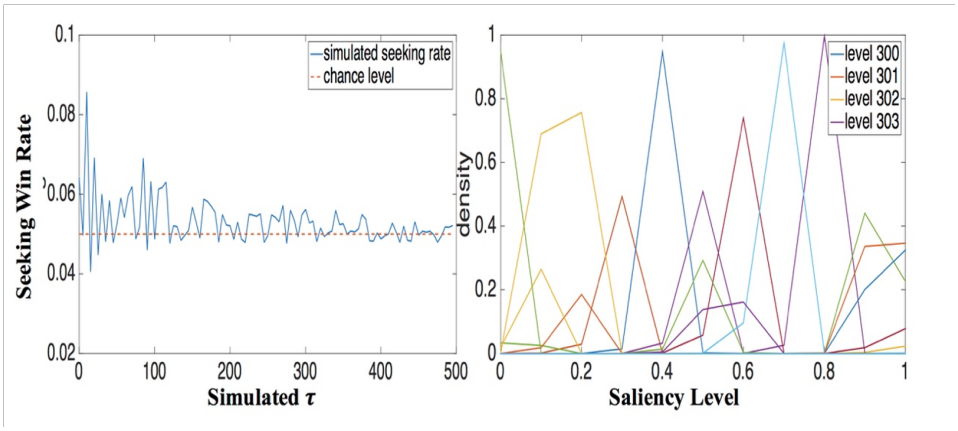

Figure 2.C.1: SCH prediction for a high level population.

The left figure simulates the seeking win rates with best-fit μ and λ with respect to different τs. The right figure shows examples for high-level strategies on the saliency domain.

group strategic level, we chose to keep one entire session that contains all subjects' data in order to not mix session data and perhaps undermine the estimation of τ.

The converging properties

A remaining question is whether seeker's advantage will disappear if people are highly strategic. Figure 2.C.1 shows the simulation results from the SCH model using the best fit μ and λ across $\tau = 0$ to 500. As can be seen, the saliency bias will only affects the seeking win rate when the level is low. The seeking win rate will converge around the equilibrium level 0.05 when the strategic level is high. In this experiment, we only observed very low level behaviours. The right figure 2.C.1 shows examples of high level strategies on the domain of saliency. Though the seeking win rate converges, the higher order behaviors actually do not converge, which means that the convergence in seeking win rate is actually due to the mixing of more types.

2.D Quantal response equilibrium prediction on saliency

Set-up

We define a game **G** with a finite strategy set 1,2....n...N. Assuming that each agent has general utility function:

$$u_n = M + \mu S_n$$

,in which n is the choice number, M stands for any monetary reward one could get, and S_n is the corresponding saliency value of choice n.

Following the notations from McKelvey and Palfrey (1995), in a game setting, the utility for player i conditional on other players using strategy p_{-i} becomes:

$$EU_n(p_i, p_{-i}) = EM_i(p_i, p_{-i}) + \mu S_n$$

M_i is the utility function from the game.

Quantal response equilibrium (QRE) solution

In general, every element p_n in a logit QRE strategy p satisfies:

$$p_n = \frac{e^{\lambda EU_n(p_i, p_{-i})}}{\sum e^{\lambda EU_m(p_i, p_{-i})}}$$

Specifically, in a two persons hide-and-seek game, p_s and p_h, strategies for seeker and hider, respectively, constitute a QRE if:

$$\begin{cases} p_{sn} = \frac{e^{\lambda EU_{sn}(p_s, p_h)}}{\sum e^{\lambda EU_m(p_s, p_h)}} \\ p_{hn} = \frac{e^{\lambda EU_{hn}(p_s, p_h)}}{\sum e^{\lambda EU_m(p_s, p_h)}} \\ \sum p_{sn} = 1 \\ \sum p_{hn} = 1 \end{cases}$$

Without loss of generality, hider will get one unit monetary reward if hided successfully and seeker will get σ unit monetary reward if successfully found the hider. Reward can be asymmetric between roles but uniform between choices. In this way, the QRE condition can be further written in (condition 1):

$$\begin{cases} p_{sn} = \frac{e^{\lambda(\sigma p_{hn} + \mu S_n)}}{\sum e^{\lambda(\sigma p_{hm} + \mu S_m)}} \\ p_{hn} = \frac{e^{\lambda(1 - p_{sn} + \mu S_n)}}{\sum e^{\lambda(1 - p_{sm} + \mu S_m)}} \end{cases}$$

Seeker's advantage condition

Seeker's advantage simply means $\sum p_{sn} p_{hn} > \frac{1}{N}$, where N is the total number of options. The easiest way to do it is through sequential inequality. Sequential inequality: for any two group of real sequences x_n and y_n:

$$x_1 < x_2 < \ldots x_n$$

and

$$y_1 < y_2 < \ldots y_n$$

Suppose $x_{\sigma(n)}$ is an arbitrary ordering of sequence x_n

$$\Rightarrow x_1 y_1 + x_2 y_2 \ldots x_n y_n \geq x_{\sigma(1)} y_1 + x_{\sigma(2)} y_2 \geq x_1 y_n + x_2 y_{n-1} \ldots x_n y_1$$

according to which, the condition, $\sum p_{sn}p_{hn} > \frac{1}{N}$, is equivalent to, for any n and m:

$$p_{sn} > p_{sm} \Rightarrow p_{hn} > p_{hm}$$

This simply means that the sequence p_s and p_h is inverse correlated. As one of them increases, the corresponding other one must decrease.

To get this, let us start from condition 1. In the first two equations, the denominator will not change with n, thus they can be viewed as constant for studying seeker's advantage. Condition 1 further becomes:

$$\begin{cases} p_{sn} & = C_s e^{\lambda(\sigma p_{hn} + \mu S_n)} \\ p_{hn} & = C_h e^{\lambda(1 - p_{sn} + \mu S_n)} \\ \sum p_{sn} & = 1 \\ \sum p_{hn} & = 1 \end{cases} \tag{2.2}$$

In which $C_s = \frac{1}{\sum e^{\lambda(\sigma p_{hm} + \mu S_m)}}$ and $C_h = \frac{1}{\sum e^{\lambda(1 - p_{sm} + \mu S_m)}}$.

Take the first two lines,

$$\frac{p_{sn}}{p_{hn}} = \frac{C_s e^{\lambda\sigma p_{hn}}}{C_h e^{\lambda(1-p_{sn})}}$$

$$\Rightarrow p_{sn} e^{\lambda(1-p_{sn})} = \frac{C_s}{C_h} p_{hn} e^{\lambda\sigma p_{hn}} \tag{2.3}$$

Equation 2.3 holds for every choices between hider's strategy and seeker's strategy. It is trivial that RHS is a monotonically increasing function in p_{hn}. In order to make LHS to be an increasing function with respect to p_{sn}, we need to take derivative with respect to p_{sn},

$$LHS = (1 - \lambda p_{sn}) e^{\lambda(1-p_{sn})}$$

The monotonicity depends on parameter λ. The condition for LHS to be an increasing function becomes:

$$p_{sn} < \frac{1}{\lambda}$$

Notice that p_{sn} is less than one. In that case, if $\lambda < 1$, this increasing condition always holds, which means as p_{sn} increases, p_{hn} also increases. According to the sequential inequality, $\sum p_{sn}p_{hn} > \frac{1}{N}$, thus, there will be a seeker's advantage. For an extreme case, when λ is approaching zero, people are choosing salient locations, which will generate a seeker's advantage because of the correlation between two roles.

When λ is very big, the increasing condition never holds since the summation needs to be 1. In that case, as p_{sn} increases, p_{hn} decreases. This would cause a hider's advantage also according to sequential inequality.

When λ is in between, there will be a cutoff point P_c. If $p_{sn} < P_c$, as p_{sn} increases, p_{hn} also increases. If $p_{sn} > P_c$, as p_{sn} increases, p_{hn} decreases. In this case, seeker's advantage varies and will has a break even point for some λ.

Chapter 3

MODELING CHOICE-PROCESS DATA ON STRATEGIC GAMES – AN APPLICATION OF HIDDEN MARKOV MODEL (HMM)

3.1 Introduction

There are two ways discussed in the literature hypothesizing how people play games: one is that everyone has a global vision about the game and knows which choices are optimal beforehand, which is often referred to as equilibrium play. In the second way, players are not that optimal. They only have bounded rationalities and make choices based on their limited cognition (Alaoui and Penta, 2015; Friedenberg, Kets, and Kneeland, 2016; Alaoui, Janezic, and Penta, 2020; Kübler and Weizsäcker, 2004; Brandenburger, Danieli, and Friedenberg, 2017). To understand the second way, the behavioral strategies of how human actually play games, most studies approach this question in a "revealed way": using choice data only, which provides very limited information. For instance, a player who does not read the rules of the game and play randomly versus another player who fully understands the game and comes up with a random strategy, we cannot distinguish between them only based on their choices. Exploring the "cognitive" process of how strategies form helps us understand human strategic plays. One way to study the cognitive part is to look directly at human brains (Coricelli and Nagel, 2009; Hampton, Bossaerts, and O'Doherty, 2008; Ong, Madlon-Kay, and Platt, 2018). A much simpler way is to look at choice–process data such as response time, gaze movements, or mouse traces, which are more feasible to get than brain data. These choice process data serve, as a hint of what people are thinking about. Though how to systematically model these data and what kind of new predictions they are able to make remain questions.

This paper presents a novel machine learning model framework, a hidden Markov model (HMM), to help better understand the underlying mental processes for game plays using choice process data, and to make new predictions. HMM is a method widely used in machine learning and cognitive science, although rarely in economics. A typical HMM consists of four parts: a hidden state set Ω, probabilities of transitions between the hidden states P, and associations between the hidden states and observable behavior B (called "data emissions"). In this paper, we will present

an application assuming that the association probability is Gaussian distributed.

In a game setting, hidden states can be regarded as mental states that cannot be observed directly. Different hidden states will have different choice generating schemas. Agents jump among different hidden states as if they were experiencing a "back and forth" process until some stopping criteria is reached.[1] What can be observed are different traces generated from these underlying hidden states, for example, mouse traces or gaze traces. These traces will be treated as probabilistic estimations of the unobserved mental states. We will use these two types of choice process data as our observable data sets (also referred to as emissions).

A natural extension of the current HMM is to include visiting durations of hidden states to the current HMM and the extended model turns into a continuous-time HMM (ctHMM). A ctHMM simply takes the time length of eye fixations into consideration so that the input of the model consists of two parts: eye-fixated strategies and fixation time spans. Staring at different strategies with shorter or longer period of time might carry information but it has not been studied yet. In the cgtHMM, besides the previous model structure, the model will further learn a stochastic process on how long a decision maker thinks about a hidden strategy. For simplicity, we will call the model without visiting durations discrete hidden markov model and the model with visiting durations continuous hidden markov model. [2]

The naive example

Let me bring up a common but intuitive scenario in games and try to explain how hidden markov models are connected with strategic processes. Consider the 2×2 matching-penny game shown in table 3.1.1. On behalf of Ann (the column player), if she thinks Bob (the row player) is playing any strategy other than pure random (equilibrium strategy), her best response will be a pure strategy, say "L". The hypothetical Bob in her mind would best respond to choose "T". She should next choose "R", and so on so forth. Such "back and forth" processes appear in many types of games and it is inevitable for most people to leave choice process signals of what they are thinking about, for example, more gaze transitions or longer response time.

Figure 3.1.1 demonstrates a hypothetical application of HMM on the matching penny game from the perspective of player Ann (table 3.1.1). Ann starts from a

[1] Either a natural stop or the time is limited.

[2] Notice that discrete in time is a different notion from discrete in hidden states. We will see applications with continuous hidden states later.

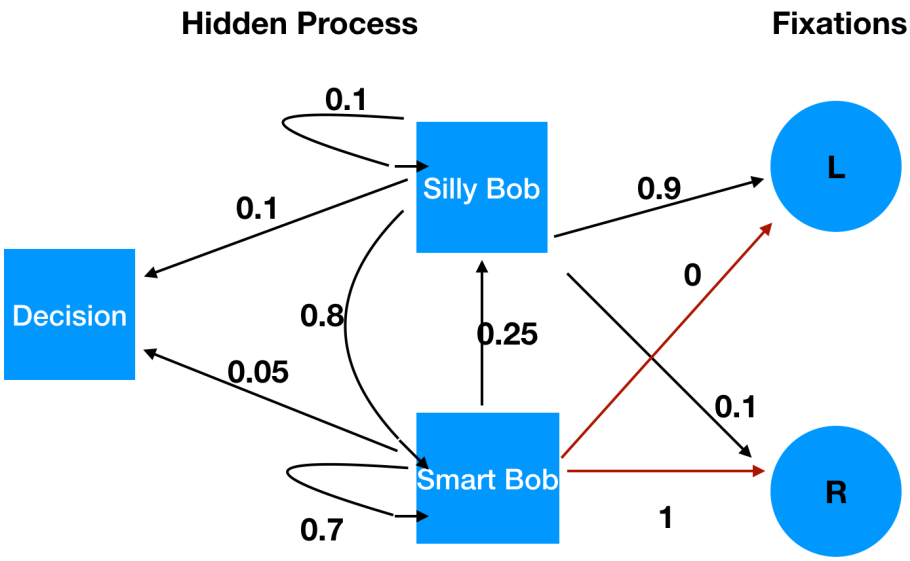

Figure 3.1.1: A hypothetical example.

	L	R
T	(0,1)	**(1,0)**
B	**(1,0)**	(0,1)

Table 3.1.1: The most simple example

prior distribution of two hidden states here: Silly Bob and Smart Bob. Ann can either choose to stay at the current state or switch to anther state, or makes a decision and exist the game. At each hidden state, she prefers different strategies and will look her current preferred strategy more. For example, at Silly Bob state, she thinks L is probably the best option so she will look at it 90% of time while at Smart Bob state, she will only look at action R. If we further consider a continuous-time case, when Ann thinks about different strategies, she will have different time distribution of when to switch to the next state. Once she decides, she will choose according to the emission probability of the last state she visited. For example, if she decides as Smart Bob, we will observe her choosing R.

What can be learned from an estimated model?

As mentioned earlier, unlike most of the decision models which require a choice data set as an input, the HMM we constructed does not need choice data at all. What the model needs are one of the two specific types of choice process data: gaze movements or mouse tracking data. In this paper, I will only focus on an application with gaze movement data. Given a data set which is collected during a game task

and an assumption of the number of hidden units, an HMM model is able to learn the best-fit transition probabilities (P) among the hidden states, an association (B) between hidden states and observed data, and a prior distribution (A) among hidden states initially. The HMM serves as a new methodology that can model choice process data in games and infer the dynamic belief formation process. Furthermore, the estimated model can be then treated as a virtual game player and generate data. In machine learning, such a model is often referred to as generative model.[3]

Continuous-time model, in addition, is capable of analyzing and predicting response time. It estimates probability transitions ($P(t)$) at each time point. Thus, this model can predict an estimated response time given a trained ctHMM model, while none of the traditional methods can do it. Other than that, ctHMM can further predict choices at any given time limit as if an agent is facing a time limit pressure. This progression is important as it enables us to manipulate the time limit and study the comparison between naive strategies and sophisticated strategies. To solve a similar problem, Agranov, Caplin, and Tergiman (2015) used a different approach. They applied a "choice process" protocol and tried to identify naive plays in a traditional guessing game.

New predictions on levels of reasoning

What currently lie between "Nash" and the reality are modern behavioral game models (Nagel, 1995; Stahl II and Wilson, 1994; Crawford and Iriberri, 2007b; Camerer, Ho, and Chong, 2004; Alaoui and Penta, 2015; Kneeland, 2015). Many of these models point to a common core idea: iterated levels of reasoning, which means players start from a naive strategy and update their belief towards the direction of the best response function. A decision will be made after several steps of such process. The number of steps is defined as "levels".

Identifying such levels is actually hard for two main reasons: 1) the system highly depends on the assumption of the starting point, level zero strategies, 2) for many cyclic games[4], the same choice can be explained using different level strategies. Previous studies avoid these two problems by using specially designed games such as 11-20 money request game, p-beauty contest game, or ring game [5] (Nagel, 1995; Arad and Rubinstein, 2012; Lindner and Sutter, 2013; Kneeland, 2015). For most

[3]Generative model is often compared with discriminative model. The latter is used to distinguish different classes while the former can be used to generate new data. HMM is a typical generative model.

[4]Exist strategy s_i, s.t $s_i \in \mathbb{B}(\mathbb{B}(s_i))$

[5]Ring game identification does not assume a level zero behavior and uses a different system.

other types of games, type identification process remains problematic. In this paper, we provide a new identification method of strategic levels based on gaze transition data and HMM method. It can be generalized to any games and does not depend on a level zero assumption. Furthermore, the identification is done on a single trial basis. We used a data set collected in saliency games as a demonstration, and will present what the result looks like. The estimated mean level is consistent with previous findings.

3.2 Discrete Model (HMM)

A discrete HMM constructed for games follows a standard framework in Murphy (2005).

Definition

A Hidden Markov Model $\mathbb{H} =< P, B, \Pi >$ consists of:

- A finite set of states $\Omega = \{h_1, h_2..h_i..h_n\}$. n is pre-specified.
- A transition matrix P, each entry p_{ij} is the probability of transitioning to state j from state i.
- A set of observation states O.
- A projection matrix B, mapping the hidden state space Ω to O. b_{mn} is the probability of observing choice n condition on current hidden state is h_m.
- A $n \times 1$ prior vector Π, indicating the initial probability distribution on h_is.

In game context, the observed gaze fixation data set consists of the observation states O. Specifically in one trial, we will observe a vector of sequenced fixations on the action space of the game, denoted by $\mathbf{f}$. As O is a probabilistic projection from Ω, hidden states can be therefore understood as hidden mixed strategies. Two states are different and are not projected to Ω. One is start state, which means a game has not been started. Another is end state, which means a decision is generated. For simplicity, we will always denote h_1 as start state and h_n as end state. A process from h_1 to h_n is the strategic reasoning process that we are interested in.

Given a data set O and game $\mathbf{G}$, a Hidden Markov Model $\mathbb{H}$ can be estimated using the Baum–Welch algorithm (a special case of the expectation maximization algorithm). For details about the estimation process, see Appendix. An estimated

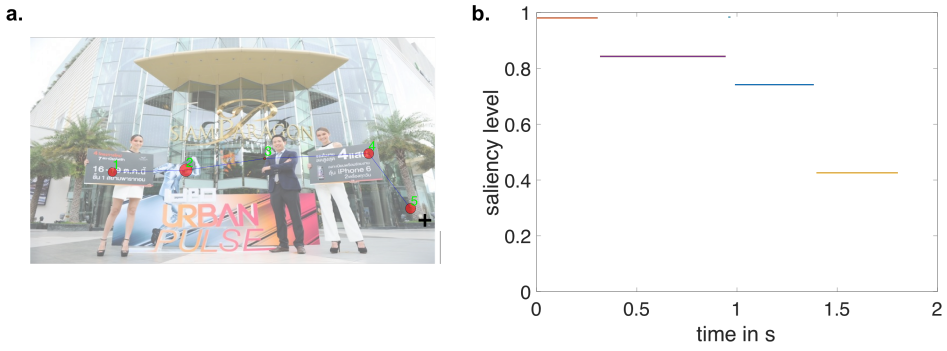

Figure 3.3.1: An example of fixation data sequence in one trial.

a) shows five fixation locations overlaid on the original image and the time order of them (labeled in green). The radius of the red circle indicates the lasting time of each fixation. b) plots the saliency levels of the fixations in figure a) against the corresponding timestamp in that trial. This specific trial starts at time zero.

HMM can be used to compare different reasoning processes under different strategic goals.

In the next section, we will demonstrate an application of HMM on saliency games, and show how an HMM behaves when facing hiding, seeking, or matching tasks.

3.3 An application of the Gaussian hidden Markov model (gHMM)

Figure 3.3.1a illustrates the sequence of five fixations for a single trial in this particular game. The order of fixations is shown by discrete numbers (the first fixation denoted by "1", the second by "2", etc.). The size of the dots at each fixation location is scaled by the length of the fixation duration.

As is well-known in human visual perception, it is evident in this example that the eye is performing a series of rapid, discrete fixations (about 100-300msec long), transitioning rapidly between fixations. Figure 3.3.1b shows the same data in a time series, which also plots (on the y-axis) the SAM-model salience of the location at which the eye fixates over time.

We start with the observation that the eyetracking indicates that fixation lengths are quite similar. Then it is reasonable to approximate the dynamic attentional process as a series of discrete transitions between states corresponding to different levels of salience (with time in each state fixed in length and therefore suppressed).[6] A

[6]Some readers will be familiar with dynamic process models which hypothesize a hidden accumulator or "drift diffusion" models (DDM) leading to choice (Ratcliff and McKoon, 2008; Ratcliff, Smith, et al., 2016). These models are designed to explain choices, error rates, response

continuous version is described in the next section, once the simpler discrete version builds intuition.

The HMM we used in saliency games hypothesizes four states. Other assumptions about hidden state numbers have also been tried and four gives the optimal BIC value (Figure 3.B.1). We will use a letter notation instead of sub-numbers to help readers better understand this example. One is a single starting state S. The two internal states, L and H, represent whether the salience level of the location being perceived is low or high (more precisely, whether salience levels are drawn from either of two Gaussian distributions of saliency levels, one with a low mean and one with a high mean). Another is an ending decision state D. A trial ends when one of the two internal states, L or H, end the trial by transitioning to D. The decision is assumed to be the fixation accompanying the last L or H salience state that transitioned to D.

The crucial unobservable parameters are the probabilities of transitions between states. The starting state effectively has two transition probabilities, $p(L|S)$ and $p(H|S)$, as well as a rare transition $p(D|S)$.

The parameters that need to be estimated from the eye tracking data are the transition probability matrix P, and the means and variances of the Gaussian L and H salience distributions. The estimation is done using the standard Baum-Welch algorithm through the HMM estimation toolbox in Matlab (Murphy, 2005) (see Appendix for details).

Figure 3.3.2 shows a graphic display of transition probabilities between states, and the estimated Gaussian distributions of salience in the L and H states. There are three figures, corresponding to matching, hiding and seeking.

In all three games, subjects transition from Start S to state H most often (over 90% of time). Remember that the HMM is trained on the actual eyetracking data. This high transition rate $P(H|S)$ is evidence that the SAM, which was trained on free gaze by other people, is also doing a good job predicting where our game-playing subjects look initially.

However, the fixation transitions afterwards vary by games. Subjects transition from H to the less-salient L node in the hider-seeker game more often than in matching game, at rates of 32% for hiders and 20% for seekers. Transitions to the decision

times (RTs), and differential RTs for correct and erroneous choices (Ratcliff, Smith, et al. (2016) or so). DDM models are not ideal because, as Figures 3.3.1ab indicate, the core mental process is transitions between salience levels rather than accumulation towards a boundary.

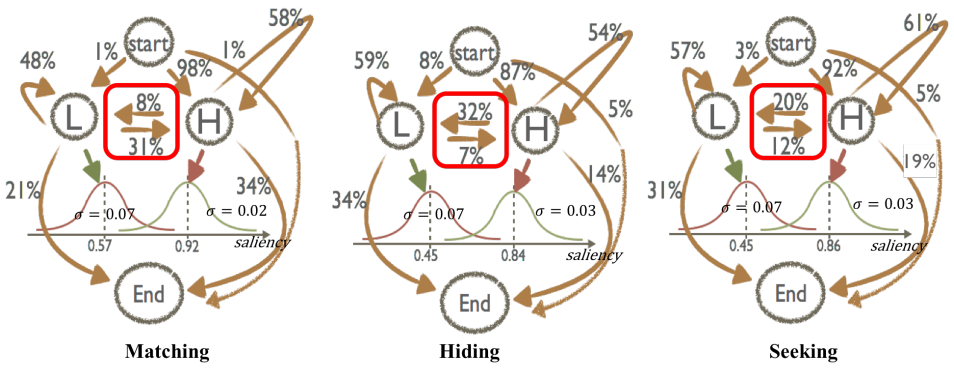

Figure 3.3.2: Estimated HMM model.

Each column in the figure shows in one game, the estimated transition matrix P among hidden states and the parameters of Gaussian transition processes.

node D are also different across games. Conditional transitions from H states are more common than from L states in matching games (34% vs 21%). Hiders are especially unlikely to conditionally transition out of salient H states to Decision (14%) compared to seekers (30%), reflecting strategic thinking.

3.4 A novel way for level identification

We argue a new strategic level identification method based on fixation transition data. We follow a simple assumption: when people are reasoning among different options, their eyes follow what they are thinking about. Under this assumption, we will be able to observe that higher level players jump back and forth between different options multiple times while low level players have much fewer fixations and make their choice faster. We can use the HMM method and fixation data to back up an estimated k level on a single trial basis given any game structure.

Given an estimated HMM model $\mathbf{M}$, suppose a vector of fixation emissions denoted by $\mathbf{f} = [f_1, f_2, ..f_m]$ like figure 3.1.1 is observed in a single trial under game $\mathbf{G}$. $\mathbf{f}$ can be projected to $\mathbf{h}$, an ordered vector in the estimated hidden state space by maximizing the posterior probability, $P(h_i|\mathbf{f}, \mathbf{M})$. Formally, a project function $g_{o\to\omega}$: $O \to \Omega$ is defined as:

$$g_{o\to\omega}(\mathbf{f}) = \underset{\mathbf{h}}{argmax} P(\mathbf{f}|\mathbf{h}, \mathbf{M})$$

Now, the fixation vectors are all transformed to hidden states vectors. In the saliency game example, a hidden state transition vector looks like $[H, S, H, H, H, D]$.

With the transformed data set, we will be able to estimate the reasoning levels. First, we assume that the first fixation is the level zero response as level zero behavior by definition comes from initial reactions. Other studies also show evidence supporting that level zero behaviors should be the salient options (Crawford and Iriberri, 2007a). It is well known in other fields like psychology and neuroscience that first fixation is highly correlated with saliency level (more than 90%). Each directional transition which fits the best response function of game **G** adds one more level. An estimated level using HMM will be the sum of all such added levels. For example, a transition vector $[S, S, H, S, S, D]$ in hider-seeker games will be identified as level 2.

Formally, given a transition vector **h** in a single trial and the best-response function **B** of the game, the strategic level is defined by:

$$\mathbf{k} = \sum_i \mathbb{1}(h_{i+1} = \mathbf{B}(h_i))$$

We apply the proposed level identification method to the fixation data set from saliency games. Figure 3.4.1 shows the estimated level distributions of three different games. These level identification results fit nicely using a Poisson distribution, which is consistent with the assumption made in the cognitive hierarchy model. The best-fitted τs are 0.5, 0.97, and 0.90 for matching, seeking, and hiding, respectively.

Similarly to previous models, our model also indicates that most people are lower than level three. What makes it different is that this new method does not have identification issue as models using choice data only are hard to identify for most types of games.

Figure 3.4.2 compares the mean of population levels by using our method and two previous methods: cognitive hierarchy model and level-k model (for details about these two benchmark models, see Chapter 1). As can be seen, the estimated mean level of the new method lies between the mean level from the CH model and the level-k model. Furthermore, using our new method, the strategic level in hiding game is not different from that in seeking games ($p = 0.24$), while in matching games, people behave as lower types than in the hider-seeker game ($p < 10^{-4}$).

3.5 Extended continuous-time model (ctHMM)

The fitted discrete HMM model in section 3.2 can be further extended to a continuous Markov model (ctHMM) and predict choices under time pressure. The ctHMM preserves all the properties of the discrete HMM, and it further contains the fixation lengths and transition time.

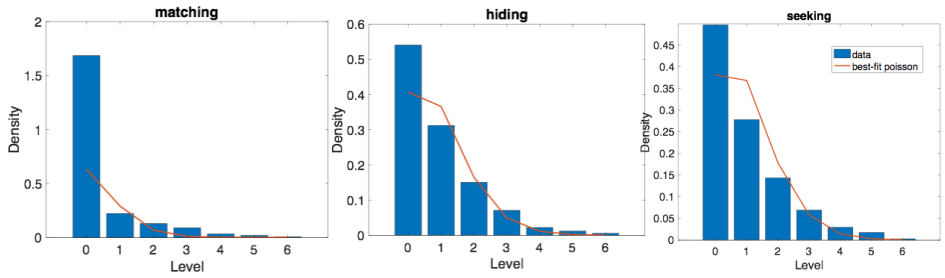

Figure 3.4.1: Level identification from gaze transitions and HMM.

The blue bar is the distribution of levels identified using gaze transition method. The red line is the best-fit Poisson distribution of the gaze identified levels. A Poisson distribution doesn't fit matching game data as well as hiding data or seeking data.

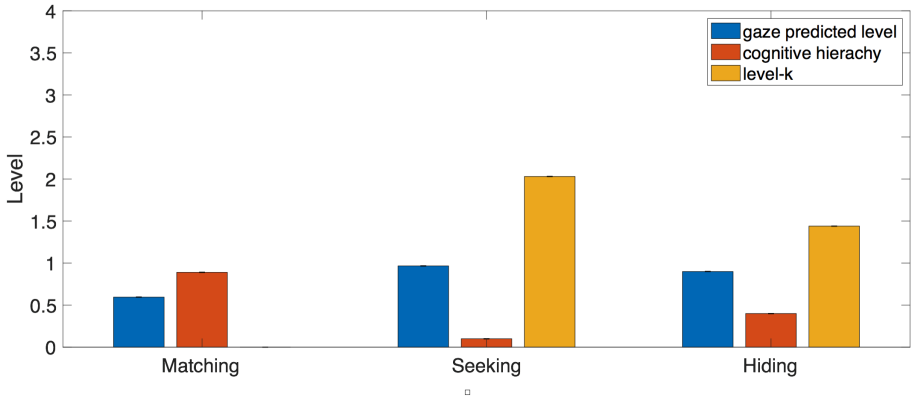

Figure 3.4.2: Comparison of levels between the gaze estimated method and other methods.

The blue bars are the average levels estimated using gaze data. Red bars are those levels using cognitive hierarchy method (also reported in the first chapter). Yellow bars are the levels using level-k method.

Formally, a ctHMM is composed of a transition probability matrix P, an initial probability vector Π, a mapping B from hidden state space Ω to data emissions in observational space: O, and the probability of transiting to state j from i at time t denoted by $P_{ij}(t)$. In Section 3.2, we already estimated the transition probability matrix $\hat{P}$ and $\hat{\mu}$. In this section, our main goal is to estimate $P_{ij}(t)$, the dynamic profile of transition probability over time.

The ctHMM will then be able to predict how long the decision maker will stay in each hidden state and the probability of observing a particular data emission at each time point beginning from time zero. A sample prediction looks like: at time 2s, with probability 0.7 the decision maker is considering hidden strategy s_i and with

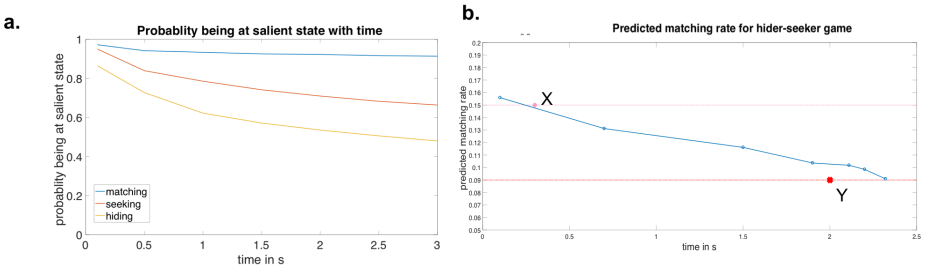

Figure 3.5.1: CtgHMM model predictions of saliency over time.

(a)The figure shows the probability of people being in the high-salience state as time goes by in the unit of seconds (b) The simulation results of ctgHMM predicting the matching rate averaged against reaction time over all the task images. Point X is the estimated average matching rate of the time-pressure group and point Y is the estimated average matching rate of the baseline (no time pressure) group.

probability 0.2 we will observe her looking at action o.

The ctHMM model makes predictions about what actions will be looked at, and chosen, if the exogenous limit on response time is put. A ctHMM can answer the following question: If the dynamic strategic process of the player is artificially truncated earlier, what does the player do? Does she not have time to find a reasonable strategy, and hence suffers from losing more in the game under time pressure?

We will demonstrate the results of such prediction in the next secion.

3.6 Prediction of time-pressure condition in games

We will also use the data set from the saliency game to validate the prediction. Similarly to the previous secion, the association function is assumed to be Gaussian, thus the model is called continuous-time Gaussian hidden markov model (ctgHMM).

The data set we have consists of what are called “emission vectors”. An emission, in ctgHMM is the saliency value of the fixation location in the discrete model. In ctgHMM, the saliency values of these fixations are combined with the time span of eye fixation at the location of interest. With the estimated model in the discrete case, we follow the standard MLE method described in Ross (2002) and a well-known cthmm package tutorial[7] and further estimate $P_{ij}(t)$ for each game (for details about the estimation, see Appendix). The result is shown in Figure 3.5.1A. Note that a maintained assumption in this exercise is that players do not adjust their HMM

[7]https://research.fhcrc.org/content/dam/stripe/etzioni/files/cthmm%20manual.pdf

structure due to time pressure. If this assumption is wrong, the predicted effects of time structure may not be accurate.

The model predicts that in all three games, people start at salient locations in the first few hundred milliseconds. In the matching game, the decision path stays at this high saliency state and only goes down a little over time. In hider-seeker games, the probability of being at the high-saliency state drops much faster than in the matching case, and drops more quickly for hiders.

If time pressure reduces thinking levels in our location-choice games, that should increase the seeker's advantage.

We measured the seeker's advantage between control (time limit = 6s) and time-pressure treatment groups (time limit = 2s) by collecting new data. In the control condition, the online group still shows a significant seeker's advantage at precisely the same rate, 0.09, as in the local lab population. This seeker's advantage increased significantly to 0.15 ($p < 10^{-4}$, t-test) in the time pressure feedback condition.[8] This result is shown in Figure 3.5.1b, labeled by point X (time-pressure matching rate) and point Y (baseline matching rate), indicating that the cgtHMM model predicts the seeking rate at two different time points reasonably well. Thus, our understanding of how saliency interacts with the depth of strategic thinking enabled us to predict a change in seeker advantage due to time pressure.

While the effect of time pressure might seem entirely obvious, keep in mind that no model in game theory actually predicts it, because there are no theories integrating salience and dynamic attention. The prediction that the seeker's win rate will go up under time pressure is certainly intuitive, but it rests entirely on the hypothesis that under time pressure, hiders cannot inhibit the automatic tendency to look at – and then choose – salient locations. Furthermore, the ctgHMM goes even further than that time pressure intuition by delivering a continuous predicted profile of seeker win rates as game durations vary. It should be easy to imagine interesting experiments testing such numerical predictions on these data and many other where dynamic attention is measured and understood (e.g., HMM transitions between lower and higher level-k thinking).

To investigate whether this increase in seeker's advantage was due to subjects choosing more salient locations under time pressure, we compared the mean saliency level

[8]The same bootstrapping method that was used in the previous seeker's advantage estimation is also used here. Note also that the seeker win rate is higher, 23%, in the no feedback condition.

		NoB	Matching rate
Nash mixed prediction			0.071
Matching game	Local control	1147	$0.64(4*10^{-4})$
	online control	458	0.63(0.001)
	online time-pressure	388	0.63(0.02)
Hider-seeker game	Local control	1090	$0.09(3*10^{-4})$
	online control	441	$0.09(6*10^{-4})$
	online time-pressure	372	0.15(0.001)

Table 3.6.1: Realized matching rate for time pressure group

of location choices in the control and time pressure treatments. Hider mean choice saliency increased under time pressure, from 0.52 to 0.59. The corresponding seeker number is from 0.61 to 0.81.

3.7 Discussion

In this paper, we establish a new machine learning method, a hidden markov model, which takes advantage of choice process data (particularly fixations) and can explain the strategic dynamics in games. This model is able to learn the best way to partition an action set into several hidden states. We present an application using the saliency game in chapter one, where the action set contains more than one million options. The model performed well with only two hidden states and successfully learned that people started with a similar first thought to consider salient choices while turned into different paths afterwards under in three different games, coordination, hiding and seeking.

Our model also contributes to the current level-k literature. We specified a new method to type classify a strategic level for each single choice using fixation data and a trained HMM. The fitting results consolidate previous findings that levels are distributed in Poisson and the average level is often less than three. The new type specification method also solved two problems which previous methods could not achieve: 1) The level identification does not rely on a special game structure and can be generalized. 2) The level is directly defined on the process itself rather than choices. It revolutionized an old concept that strategic levels originate from individual individual differences. Instead, strategic levels come from a combination of performance in a particular game and the average performance of a person.

Last but not least, we build a continuous-time HMM on the top of the discrete one. The continuous model is able to process fixation time lengths and use them for decision prediction at continuous time domains. In particular, this model success-

fully predicts the game outcome of saliency game, when people are facing a time pressure.

One limitation of our study is that we do not know yet what percentage of people reason without looking. However, such limitation on HMM has less effect on the general interpretation of HMM as they are probabilistic estimations anyway, but it will underestimate the levels because we are not able to identify those higher levels who reason mentally using the gaze method.

An important open question for future research is whether this method can address the portability issue illustrated in Hargreaves Heap, Rojo-Arjona, and Sugden (2014). The new definition solves the problem to some extent, but not directly. Since HMM levels are defined on single trials instead of people, which incorporate the game structure, it naturally relaxes the condition that the same individual should be consistent with their levels. It is still meaningful though to ask this question: do people have different abilities in strategic games? If so, how should we define it? A possible future research could be to collect large data sets across subjects on different games with their gaze data recorded and to look at individual differences from an HMM perspective.

References

Agranov, Marina, Andrew Caplin, and Chloe Tergiman (2015). "Naive play and the process of choice in guessing games". In: *Journal of the Economic Science Association* 1(2), pp. 146–157.

Alaoui, Larbi, Katharina A Janezic, and Antonio Penta (2020). "Reasoning about others' reasoning". In: *Journal of Economic Theory* 189, p. 105091.

Alaoui, Larbi and Antonio Penta (2015). "Endogenous depth of reasoning". In: *The Review of Economic Studies* 83(4), pp. 1297–1333.

Arad, Ayala and Ariel Rubinstein (2012). "The 11-20 money request game: A level-k reasoning study". In: *American Economic Review* 102(7), pp. 3561–73.

Brandenburger, Adam, Alex Danieli, and Amanda Friedenberg (2017). *How many levels do players reason? An observational challenge and solution*. Tech. rep. Working paper.

Camerer, Colin F., TeckHua Ho, and Juin–Kuan Chong (2004). "A cognitive hierarchy model of games". In: *The Quarterly Journal of Economics* 119(3), pp. 861–898.

Coricelli, Giorgio and Rosemarie Nagel (2009). “Neural correlates of depth of strategic reasoning in medial prefrontal cortex”. In: *Proceedings of the National Academy of Sciences* 106(23), pp. 9163–9168.

Crawford, Vincent P. and Nagore Iriberri (2007a). “Fatal attraction: Salience, naivete, and sophistication in experimental" hide-and-seek" games”. In: *American Economic Review* 97(5), pp. 1731–1750.

Crawford, Vincent P. and Nagore Iriberri (2007b). “Level–k Auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private–value auctions?” In: *Econometrica* 75(6), pp. 1721–1770.

Friedenberg, Amanda, Willemien Kets, and Terri Kneeland (2016). *Bounded reasoning: Rationality or cognition*. Tech. rep. Mimeo.

Hampton, Alan N, Peter Bossaerts, and John P O’Doherty (2008). “Neural correlates of mentalizing-related computations during strategic interactions in humans”. In: *Proceedings of the National Academy of Sciences* 105(18), pp. 6741–6746.

Hargreaves Heap, Shaun, David Rojo-Arjona, and Robert Sugden (2014). “How Portable Is Level–0 Behavior? A Test of Level–k Theory in Games With Non–Neutral Frames”. In: *Econometrica* 82(3), pp. 1133–1151.

Kneeland, Terri (2015). “Identifying higher–order rationality”. In: *Econometrica* 83(5), pp. 2065–2079.

Kübler, Dorothea and Georg Weizsäcker (2004). “Limited depth of reasoning and failure of cascade formation in the laboratory”. In: *The Review of Economic Studies* 71(2), pp. 425–441.

Lindner, Florian and Matthias Sutter (2013). “Level-k reasoning and time pressure in the 11–20 money request game”. In: *Economics Letters* 120(3), pp. 542–545.

Murphy, Kevin (2005). *Hidden Markov Model (HMM) toolbox for Matlab Retrieved May 25 2012*.

Nagel, Rosemarie (1995). “Unraveling in guessing games: An experimental study”. In: *The American Economic Review* 85(5), pp. 1313–1326.

Ong, Wei Song, Seth Madlon-Kay, and Michael L Platt (2018). “Neuronal mechanisms of strategic cooperation”. In: *bioRxiv*, p. 500850.

Ratcliff, Roger and Gail McKoon (2008). “The diffusion decision model: theory and data for two-choice decision tasks”. In: *Neural Computation* 20(4), pp. 873–922.

Ratcliff, Roger, Philip L Smith, et al. (2016). “Diffusion decision model: Current issues and history”. In: *Trends in Cognitive Sciences* 20(4), pp. 260–281.

Stahl II, Dale O. and Paul W. Wilson (1994). “Experimental evidence on players’ models of other players”. In: *Journal of Economic Behavior & Organization* 25(3), pp. 309–327.

APPENDIX

3.A Eye tracking procedure

Subjects sat comfortably about 50cm in front of the screen while their eye movements were monitored with a Tobii TX300 eye-tracker with a sampling rate of 300Hz. Presentation of the game used the Matlab Psychophysics tool box. A nine-point calibration was performed in the beginning of the experiment. Before each sub-block, there was a drift correction phase in which subjects needed to focus on a dot in the center of the screen. The experiment only proceeded, after drift correction, if the gaze detection was still accurate. Otherwise, the eye-tracker would automatically enter a second calibration session, then continue. Between presentations of sequential images, a fixation cross was presented on a gray screen for one second to wash out any remaining gaze trace from the last image.

3.B HMM model estimation

HMM model estimation

The HMM package that we used follows the steps in the Baum-Welch algorithm in Murphy (2005) and Baum et al. (1970).

Each fixation vector, like the specific example in figure 3.3.1, constitutes one data observation. Our goal is to estimate the hidden transition matrix P and the Gaussian transition structure (G). G consists of a set of μ_i and σ_i for each hidden state. Hidden states are denoted by $_t \in \{1...K\}$. The whole parameter structure is denoted by $\theta = (\pi, P, O)$. The probability of the starting location π is known in our case, $\pi(i) = p(h_1 = i)$. $P(i, j)$ is defined by $p(h_t = j|h_{t-1} = i)$. . The set of observations O are denoted by x_t.

- Initialize a random θ to begin.

- E-step: E-step is to calculate the data log likelihood given by

$$Q(\theta, \theta_{old}) = \sum_{k=1}^{K} E[N_k^1] \log \pi_k +$$
$$\sum_{j=1}^{K}\sum_{k=1}^{K} E[N_{jk}] \log P_{jk} + \sum_{i=1}^{N}\sum_{t=1}^{T_i}\sum_{k=1}^{K} p(h_t = k|x_i, \theta_{old}) \log p(x_{i,t}|O)$$

where:

$$E[N_k^1] = \sum_{i=1}^{N} p(z_t = k|x_i, \theta_{old})$$

$$E[N_{jk}] = \sum_{i=1}^{N} \sum_{t=2}^{T_i} p(h_{i,t-1} = j, h_{i,t} = k|x_i, \theta_{old})$$

$$E[N_j] = \sum_{i=1}^{N} \sum_{t=1}^{T_i} p(h_{i,t} = j|x_i, \theta_{old})$$

- M-step: for P, we normalize the expected counts:

$$\hat{A}_{jk} = \frac{E[N_{jk}]}{\sum_{k'} E[N_{jk'}]}$$

For G:

$$\hat{\mu_k} = \frac{E[\bar{x}_k]}{E[N_k]}$$

$$\hat{\Sigma_k} = \frac{E[\bar{x}\bar{x}_k^T] - E[N_k]\hat{\mu_k}\hat{\mu_k^T}}{E[N_k]}$$

- repeat step E and step M until the resulting parameter values converge (at a convergence threshold of 10^{-4}). That is, iteration proceeds until θ is close to θ_{old}.

Simulating HMM

With the estimated continuous time model, we can predict the strategic choice for our task when people face time pressure, as shown in figure 3.5.1.

We simulate the model beginning with the starting state. At each time t, there are particles (decision makers) at the high state, low states or who already exited the game (i.e., are in the decision state). For those who arrive at the decision state, we take the high or low states they experienced before decision state as an estimation of the decision. For others, who are either on the L states or H states, we take the current state they are at as their choice saliency level.

When we add exogeneous time pressure at time t, the simulation result will be a mixture of people who naturally finished their thinking by t (those who arrived at the decision state before time t) and people who did not arrive at the decision state D in time. Our calculation assumes that those people whose HMM transitions are artificially interrupted at time t choose the saliency level of the last state, L or H,

they are in. This is quite a bold assumption, as the entire HMM – which estimated using only data from the original 6-second baseline condition – could easily change when subjects consciously know they have only 2 seconds.

After knowing the probability of choosing the low- or high-saliency point, simulating the matching rate is easy. For each simulated trial, we already know the probability of choosing at each salient state, denoted by $p(\hat{H})$. Then, we can simulate one sample, an exact saliency value with probability $p(\hat{H})$ from the estimated Gaussian H and with probability $1 - p(\hat{H})$ from Gaussian L.

This method gives a simulated choice saliency value $\hat{s_i}$ for each trial. Then we go back to the task images in the experiment and let the HMM make an artificial choice based on the interrupted saliency value. For each image, there are a set of points Q that have saliency level approximately equal to the $\hat{s_i}$ selected by the algorithm. We randomly choose one point in this set on each simulated trials. This is done for two players independently, and the matching rate of the two points selected with saliency $\hat{s_i}$ is computed.

	Matching	Hiding	Seeking
μ_H	(0.54,0.60)	(0.42,0.48)	(0.42,0.48)
μ_L	(0.90,0.91)	(0.83,0.86)	(0.85,0.87)
$P(H\|Begin)$	(0.97,0.99)	(0.84,0.89)	(0.90,0.93)
$P(L\|Begin)$	(0.01,0.02)	(0.06,0.11)	(0.02,0.04)
$P(End\|Begin)$	(0.004,0.01)	(0.04,0.06)	(0.04,0.07)
$P(L\|L)$	(0.56,0.61)	(0.50,0.57)	(0.58,0.63)
$P(End\|H)$	(0.32,0.35)	(0.12,0.16)	(0.17,0.19)
$P(L\|H)$	(0.07,0.10)	(0.28,0.36)	(0.19,0.24)
$P(H\|L)$	(0.24,0.39)	(0.05,0.10)	(0.09,0.14)
$P(L\|L)$	(0.42,0.55)	(0.56,0.61)	(0.54,0.60)
$P(End\|L)$	(0.17,0.25)	(0.32,0.36)	(0.29,0.34)

Table 3.B.1: Bootstrapped confidence interval for estimated HMM parameters

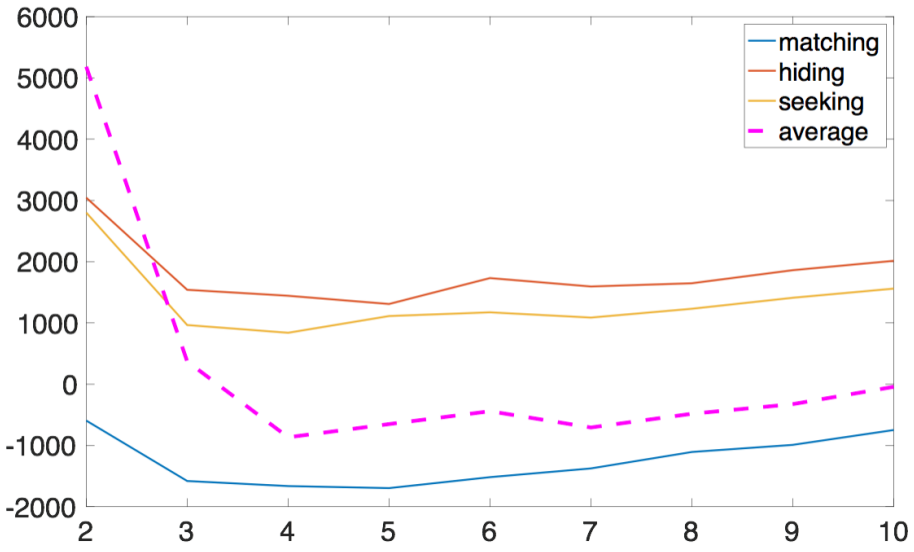

Figure 3.B.1: BIC values of different hidden units.

The figure shows how BIC values change with respect to different numbers of hidden units. The dashed line is the BIC values of a combining model with all three games.

Chapter 4

MANIPULATING VISUAL SALIENCY IN SIMPLE CHOICE PROBLEM

4.1 Introduction

People encounter decision makings when multiple choices are presented together, and there exist "the right answers" among those options. But finding these right answers may be challenging and costly. For example, merchants like to bundle various products on holidays as gift packages, and you want to buy a bundle with the greatest monetary values. Even if each item's market price in each package is known, but one may still end up with a package that does not save her any money because finding the optimal decision is costly.

Various reasons may contribute to the failure of finding the right answers. Perceptual bias is one of the most important factors among them, leading to uneven attention distributions during the decision process. As human beings rely heavily on visual information as the input information to make decisions, a property of the decision environment, visual saliency, which causes perceptual biases in the human visual system, becomes particularly important. In this chapter, we will measure and manipulate the visual saliency property of a natural and complex decision environment and show how visual saliency affects multi-items simple choice problems with objectively correct answers.

We designed a novel experiment where choices are sets of fruits which are assigned an explicit economic value. Sets of fruits are presented visually, and SAM-predicted salience and induced value are either opposed (incongruent) or similar (congruent). This paradigm enables a measurement of the differential influence of salience, independent of value, on choice. Keep in mind that the scientific risk here is that salience is entirely predicted by the parameter-free SAM model, and by how we arrange objects visually. Top-down valuation could easily wipe out transitory (1-5 sec) effects of visual salience that the later attention pattern will be drawn mostly by the unit value of the fruits (We have a long training session for subjects to learn and get familiarized with the fruit value).

We need new models

This possible attention effect that we talked about above can hardly fit in either framework of two major models in decision theory: rationality and bounded rationality. Both models assume a passive attention process that the attention distribution is only driven by all the information needed to make a decision. The context of choices that cause perceptual biases like what an alternative looks like or whether there will be a background noise is not considered necessary to make such a decision.

For example, two real-life decision problems should be treated as different decision problems, but they are mathematically equivalent in most decision theory framework: 1) choosing between left and right; 2) choosing between letters: A and B. Choosing left or A in either problem will give a prize worth value one, choosing right or B will otherwise lead to no prize and value zero. Traditional decision models will not take into account the context of the action set, and therefore, these two problems with the same payoff structure are equivalent. But if we start to think about the context, things become different. For example, there might be many reasons like heuristics to make one believe that letter A has a larger probability to lead to a prize. Considering the belief bias caused by the context of a choice will make the theory perceptual bias compatible.

In this paper, we set up a saliency-refined rational inattention model where decision makers (DMs) get some information for free from nature through saliency, which will affect their final policies. Recall that the solution concept in the original rational inattention model can be characterized in two equivalent ways (Sims, 2003; Sims, 2006; Caplin and Dean, 2013; Caplin and Dean, 2015; Caplin, Dean, and Leahy, 2019): one is to compute the optimal strategy directly: $P(a|\omega)$ under each state $\omega \in \Omega$; another is to use the posterior-based approach where the DM first chooses an attention strategy (a signal mechanism) and then makes the decisions based on their posterior belief based on the revived signals under each state. Any solution written in the first format induces its corresponding attention strategy and posterior belief(Caplin and Dean, 2013; Kamenica and Gentzkow, 2011). The implied posterior belief is in the format of:

$$\gamma_a(\omega) = \frac{P(a|\omega)\mu(\omega)}{P(a)}$$

where $\gamma_a(\omega)$ is the probability of being at state ω given choice a. The posterior way provides a better interpretation of the information aspect where the information cost

comes from forming the attention strategy.[1]

In our model, we assume that saliency does a similar job to the attention strategy. It serves as a signal device, and the DMs first form posterior beliefs based on the saliency information. But it is also fundamentally different from an attention strategy: 1) Forming such a signal process is not part of the optimization process. We can understand it as our brain or nature pre-specified it for us so that it is not guaranteed to be optimal for solving a particular problem; 2) Though it is not optimal, it is free! The decision-maker needs to pay zero cost for this part of the information.

Since such a signal process is not optimal, after forming the saliency induced posterior, the rest of the saliency refined rational inattention model should be solved just like before, replacing the prior probability of being at each state by posterior probability other than priors.

Therefore under the new saliency-refined rational inattention model, one decision problem in the traditional definition can have different solutions given different saliency distribution. (For more details see Section 4.2)

Background

Over the past decades, the phrase "mental effort" has been mentioned more frequently in various fields, including behavioral economics, decision neuroscience, and psychology. The most common way to study "mental effort" is through nudging. By either changing decision makers' total mental energies or increasing the processing cost needed for experimental stimulus, we expect to observe behavioral shifts. It is easier to quantify behavior shifts by designing specific tasks while justifying and quantifying that certain manipulation on the decision environments towards the hypothesized direction is hard. What most people do is to propose one manipulation using heuristics.

The most trivial metric for measuring attention attendance is through self-report. Scarpa et al. (2013) found that self-reported non-attentiveness measure is meaningful and coherent with the inference measurements of awareness based on the choice data. Clearly, this type of research focuses on the overall attentiveness level.

There is also extensive literature on decision neuroscience and psychology studying how attention biases value-based decision making, mostly through experiments, and

[1] It is the Shannon mutual information between the signal and the prior.

requires eye-tracking techniques (Milosavljevic et al., 2012; Towal, Mormann, and Koch, 2013; Shimojo et al., 2003; Orquin and Loose, 2013). Many evidences show that choices not only depend on independent valuations of the products, as stated in traditional economic theory. Saliency also plays an important role in determining the final choices. Both Towal, Mormann, and Koch (2013) and Krajbich, Armel, and Rangel (2010) demonstrate the benefits of accumulator model classes in studying attention problems. Towal, Mormann, and Koch (2013) finds that a model takes into account both perceptual properties and values can predict choices better than only using values. Milosavljevic et al. (2012), instead, focus more on how to manipulate attention patterns and shows that simple choices can be manipulated by changing one single visual feature of an item, like brightness. All of these three studies rely heavily on eye-tracking techniques.

Our study is definitely closely related to the line of works like Milosavljevic et al. (2012). however, the experiment used in this paper differs from previous studies in two major aspects: 1) Previous work, including Milosavljevic et al. (2012) and Krajbich, Armel, and Rangel (2010) use subjective ratings for each individual product as measurements for underlying values. Such measurements may be noisy and inaccurate. Instead, we use a set of predetermined monetary values to have better controls over the valuations. 2) Milosavljevic et al. (2012) manipulated saliency in only one simple feature: brightness. We are trying something more ambitious: to compute the saliency directly from different stimuli. It looks just like making choices in a natural scenario where saliency distribution among choices changes as many visual features change simultaneously. This paper starts by introducing a saliency-refined rational inattention model (section 4.2) where we will show how saliency information can fit in a rational inattention model. This new model predicts that saliency helps people choose the right decision when a more salient alternative is also the more rewarding alternative, while saliency misleads people to choose the less rewarding option when a more salient alternative is less rewarding. Such saliency biases can only be observed when the marginal value of making the right decision is larger, and the learning cost is small.

In section 4.3, we will introduce our fruit decision experiment where the value of each alternative is predetermined, and saliency is measured by a well-performed algorithm: saliency attentive map. We show that the result of the experiment is well aligned with the prediction from the srRI model, where people do make more mistakes when the salient item is the less rewarding item. This effect remains robust

and even stronger when we make the test environment much harder.

Other than the choice biases, we also see a response-time effect on saliency. People spend on average more time thinking when the saliency distribution contradicts with the reward distribution. Such a result from response time suggests that saliency effects could be due to a subconscious process that cannot be compensated by simply putting more time into thinking.

4.2 Saliency in Rational Inattention Model

Following the framework of rational inattention models in Sims (2003); Caplin and Dean (2013); Caplin and Dean (2015), we will show how a rational inattention model with Shannon's entropy explains and predicts the saliency effect for our experiment. The result can be easily generalized to any two-alternative simple choice problem.

- We define a product-based state space $\Omega \times \mathbb{S}$, where Ω is the traditional state space, under which the payoff structure of actions will be different. $\mathbb{S}$ is the saliency space which is observable.
- Prior beliefs on different states are denoted as μ, where $\mu \in \Gamma$ and $\Gamma = \Delta(\Omega \times \mathbb{S})$.[2] μ_i^j denotes the prior probability for being at state (ω_i, s_j).
- $\mathbb{A}$: the action set.
- Utility function $u : \Omega \times \mathbb{A} \to \mathbb{R}$.
- A decision problem consists of $(\mu \times A)$, where $A \subseteq \mathbb{A}$.

For simplicity, we assume that an individual will gain utility v if she chooses a more valuable item and 0 otherwise.

In our case, we will only consider a two-states case for Ω, where, without loss of generality, under ω_1, item one is more rewarding, and under state ω_2, item two is more rewarding. $\mathbb{S}$ is a space that contains perceptual information about the context of each choice representative. $\mathbb{S}$ simply defines saliency. Since our experiment includes only decision environments where there are big saliency differences, we further reduce $\mathbb{S}$ to only contain two elements, one is that choice A is more salient than choice B, and the other is that choice B is more salient than choice A. Therefore, this two-by-two state space will be as Table 4.2.1 shows.

[2] Γ contains all the probability distribution over all possible states.

(ω_1,s_1)	Option 1 is more rewarding and more salient
(ω_1,s_2)	Option 1 is more rewarding but less salient
(ω_2,s_2)	Option 2 is more rewarding and more salient
(ω_2,s_1)	Option 2 is more rewarding but less salient

Table 4.2.1: State Dependent Payoff Matrix.

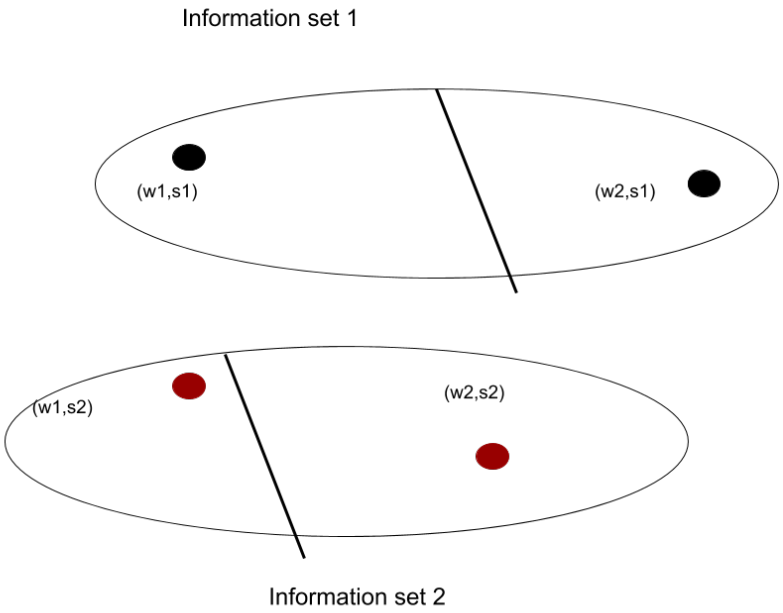

This figure illustrates that by observing different saliency distribution. A DM will either be at information set 1 or information set 2. Within each information set, it favors different payoff states.

Figure 4.2.1: Different information sets induced by saliency.

Knowing the current saliency state s cannot directly lead to the right decision as it is unrelated to the payoff function, but it serves as a signal device for implying Ω, as saliency can be easily observed by the decision maker. Therefore, DM forms a posterior belief about the current state in Ω:

$$\delta_{s=s_n}(\omega = \omega_i) = \frac{\mu_i^n}{\sum_j \mu_j^n} \tag{4.1}$$

This saliency signal process is different from the posterior approach and attention strategy denoted in Caplin and Dean (2013). One difference is that this process cannot be chosen by the decision maker but given by nature instead. Another is that the DM does not need to pay any information cost at this step for updating his belief.

With posterior $\delta_{s_n}(\omega_i)$, this problem becomes no different from the standard model. We will aim at solving the best strategy directly instead of using the posterior approach. Note that the decision maker still needs to pay the learning cost at this step because she does not have a chance to choose the optimal attention strategy through saliency. One can think of it as saliency gives part of the information for free but not all of it.

Consistent with previous models, the value of a strategy $P : \Omega \times \mathbb{S} \rightarrow \Delta(A)$ is given by the expected value of the actions, minus information costs. The decision maker now needs to choose $P : \Omega \times \mathbb{S} \rightarrow \Delta(A)$ in order to maximize:

$$\sum_{\omega \in \Omega} \delta_s(\omega) \left(\sum_{a \in A} P(a|\omega, s) u(a, \omega) \right) \tag{4.2}$$
$$-\lambda \left[\sum_{\omega \in \Omega} \delta_s(\omega) (\sum_{a \in A} P(a|\omega, s) lnP(a|\omega, s)) - \sum_{a \in A} P(a|s) lnP(a|s) \right]$$

where:

$$P(a|s) = \sum_{\omega \in \Omega} \delta_s(\omega) P(a|\omega, s)$$

In our simple case, the optimization problem shown in 4.2 becomes choosing $P(A_1|\omega_1, s)$ and $P(A_1|\omega_2, s)$ to optimize :

$$\begin{aligned} f : \delta_s(\omega_1) \, & [P(A_1|\omega_1, s) u(1,1) + P(A_2|\omega_1, s) u(2,1)] \\ & + \delta_s(\omega_2) \, [P(A_1|\omega_2, s) u(1,2) + P(A_2|\omega_2, s) u(2,2)] \\ & - \lambda [\delta_s(\omega_1) \{ (P(A_1|\omega_1, s) lnP(A_1|\omega_1, s) + P(A_2|\omega_1, s) lnP(A_2|\omega_1, s) \} + \\ & \delta_s(\omega_2) \{ (P(A_1|\omega_2, s) lnP(A_1|\omega_2, s) + P(A_2|\omega_2, s) lnP(A_2|\omega_2, s) \}] \\ & + \lambda [ln[\delta_s(\omega_1) P(A_1|\omega_1, s) + \delta_s(\omega_2) P(A_1|\omega_2, s)] [\delta_s(\omega_1) P(A_1|\omega_1, s) \\ & + \delta_s(\omega_2) P(A_1|\omega_2, s)]] \\ & + \lambda [ln[\delta_s(\omega_1) P(A_2|\omega_1, s) + \delta_s(\omega_2) P(A_2|\omega_2, s)] [\delta_s(\omega_1) P(A_2|\omega_1, s) \\ & + \delta_s(\omega_2) P(A_2|\omega_2, s)]] \end{aligned} \tag{4.3}$$

$$\tag{4.4}$$

When learning does not cost at all ($\lambda = 0$), f is linear, so that the decision maker fully distinguishes the states and chooses $P(A_1|\omega_1, s) = 1$ and $P(A_2|\omega_2, s) = 1$, implying no mistakes at all. When λ is strictly positive, we need to solve for the first order condition in order to optimize the equation. After replacing all the utilities as v and 0 according to Table 4.2.1 and using the condition all the probabilities sum

up to one, the partial derivatives of Equation 4.3 with respect to $P(A_1|\omega_1, s)$ and $P(A_2|\omega_2, s)$ become:

$$\begin{aligned}\frac{\partial f}{\partial P(A_1|\omega_1,s)} &= \delta_s(\omega_1)v - \lambda\delta_s(\omega_1)ln(\frac{P(A_1|\omega_1,s)}{1-P(A_1|\omega_1,s)}) \\ &\quad - \lambda\delta_s(\omega_1)ln\frac{\delta_s(\omega_1)P(A_1|\omega_1,s)\delta_s(\omega_2)P(A_2|\omega_2,s)}{1-\delta_s(\omega_1)P(A_1|\omega_1,s)-\delta_s(\omega_2)P(A_2|\omega_2,s)} \\ \frac{\partial f}{\partial P(A_2|\omega_2,s)} &= \delta_s(\omega_2)v + \lambda\delta_s(\omega_2)ln(\frac{P(A_2|\omega_2,s)}{1-P(A_2|\omega_2,s)}) \\ &\quad + \lambda\delta_s(\omega_2)ln\frac{\delta_s(\omega_1)P(A_1|\omega_1,s)+\delta_s(\omega_2)P(A_2|\omega_2,s)}{1-\delta_s(\omega_1)P(A_1|\omega_1,s)-\delta_s(\omega_2)P(A_2|\omega_2,s)}\end{aligned} \tag{4.5}$$

The special case of Equation 4.5 is when v/λ is zero. When v/λ becomes zero, any state independent strategy will be optimal. The good news is that this condition should rarely happen in reality because it requires either a zero value for a correct decision or an infinitely large learning cost. Besides the special case, solving 4.5 gives out the optimal $P(A_1|\omega_1, s)$ and $P(A_2|\omega_2, s)$, which are also the probabilities of people choosing correctly under each state[3]:

$$P^*(A_1|\omega_1,s) = \begin{cases} 0 & \delta_s(\omega_1) < \frac{1}{1+e^{\frac{v}{\lambda}}} \\ \frac{\delta_s(\omega_1)+\frac{\delta_s(\omega_1)-1}{e^{\frac{v}{\lambda}}}}{\delta_s(\omega_1)-\frac{\delta_s(\omega_1)}{e^{\frac{2v}{\lambda}}}} & \frac{1}{1+e^{\frac{v}{\lambda}}} \le \delta_s(\omega_1) \le 1-\frac{1}{1+e^{\frac{v}{\lambda}}} \\ 1 & \delta_s(\omega_1) > 1-\frac{1}{1+e^{\frac{v}{\lambda}}} \end{cases} \tag{4.6}$$

What we can learn from the result of Equation 4.6 is that saliency affects decision through distorting people's posterior belief on different payoff states. Figure 4.2.1 showed the simulation results of Equation 4.6 under different parameters (v/λ).

Under ω_1, option one will be the right choice to choose from. We can see that the correction rate changes when people have a biased belief that the rewarding option is option 1. The bending curves in figure 4.2.1 showed how the correction rate changes as the posterior belief of being at state one changes. The correction rate will not be affected much by the posterior belief if the value is large enough to make the correct choice or the learning cost is small. While on the other side, moving the posterior belief away from a fair one, 0.5, will change the correction rate a lot. This

[3]Without loss of generality, we will just show that $P^*(A_1|\omega_1, s)$ and $P^*(A_2|\omega_2, s)$ will be in the same format.

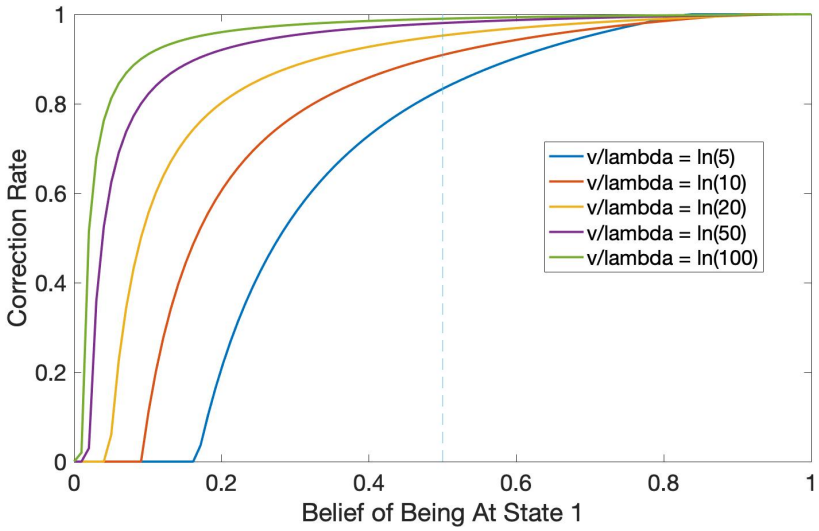

Figure 4.2.3: Correct rate simulation on saliency distorted posterior belief

is exactly what happens when the saliency bias gets involved. Such saliency biases will enhance the correct rates when the salient item is also the rewarding item and hurt the correction rate when they are opposite.

4.3 Experiment Design and Data

Task

Subjects are asked to choose one of the two fruit piles presented on the left and right part of an image, as shown in figure 4.3.1. Each fruit type unit has a unique, predetermined monetary value. The subjects become familiarized with the assigned value before the main task session. The total value of a fruit pile is the sum of the values of all fruits in that pile. The everyday analogy to this task is a retail vendor who is buying fruits at a wholesale market and has in mind a retail price for each fruit. While the vendor should be computing the resale value, the visually salient properties of the fruit (color, contrast, whether they 'pop out' from a fruit basket) are active and independently controllable by us. We can therefore construct choice sets in which visual salience (controlled by our displays) and fruit value (also controlled) are positively or negatively correlated. Subjects might ignore visual salience or might be unable to inhibit their unconscious bottom-up influence.

The main experiment is done under time limitations (10s). Subjects will be rewarded with money following an incentive-compatible design. There is a 10% chance that

the subject will get an equivalent amount of money from a randomly selected choice among all twenty rounds.

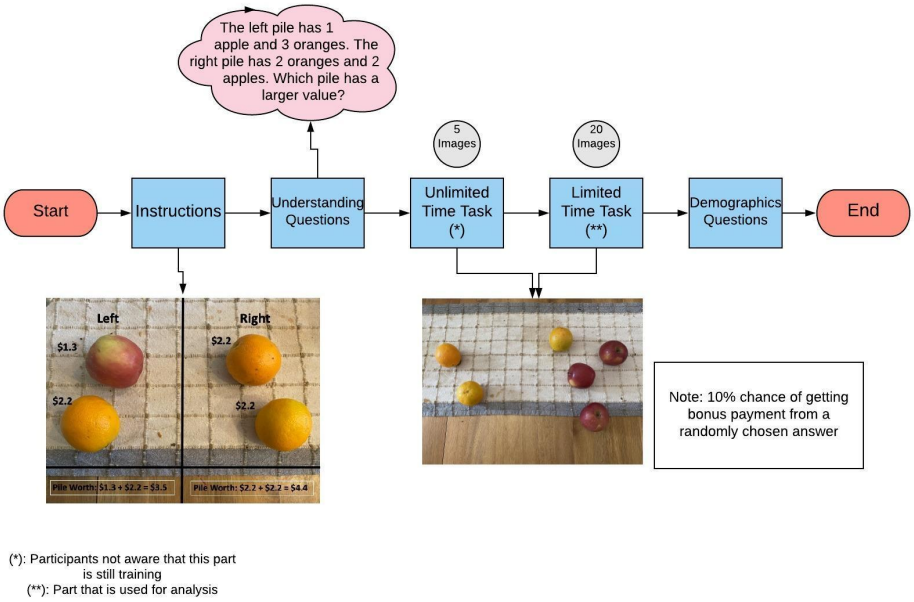

Subjects first experienced an introduction session explaining the basic tasks followed by a testing session asking questions about the rules. They then experienced an unlimited time session with N=5 images which will count into payment but only for training purposes. Afterwards, they will do a session with time limit. The experiment ends with demographic questions and payment realizations.

Figure 4.3.1: Experiment flow.

Data

N = 75 people participated this study on Prolific, a European online data collection platform following a pre-registration process on the open science foundation website (OSF). All the participants were pre-screened to have a prior approval rate of at least 70% based on their previous participation. Each subject was only allowed to participate once for all types of batches (including pilot studies). Participation from mobiles and tablets was not allowed in order to control for attention effects.

4.4 Set-up

Visual Saliency Properties of Stimuli

We used the same deep-learning based saliency model, "saliency attentive model (SAM)" (Cornia et al., 2016) [4] SAM takes the visual scene of one round (one image) and predicts where most people will allocate their attention. Examples of the outputs from SAM are shown in Figure 4.4.1bd where darker red means the location is more salient. The most salient location has the saliency value of one, and each image may have multiple local maxima at value one.

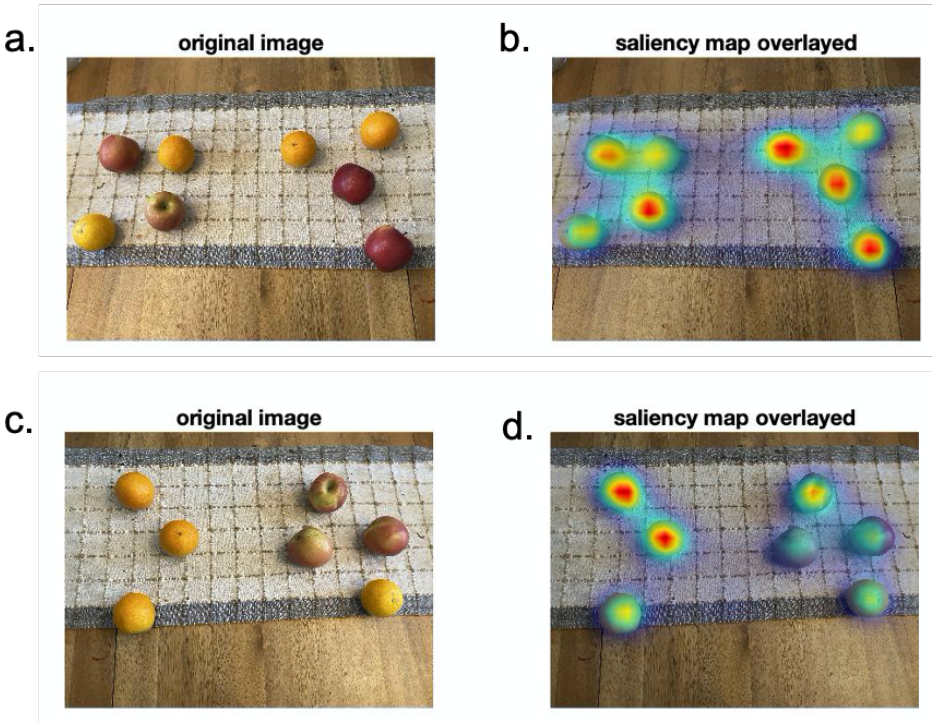

a) a sample image in the image pool with saliency centers on both sides. b) saliency prediction of image a). c) with saliency center only on one side . d) saliency prediction of c). Saliency algorithm outputs a normalized result and the most salient pixel will have saliency value one. The saliency value in the dark red area is approaching the maximum value one centered at the most salient spot.

Figure 4.4.1: Two types of saliency distribution.

Stimuli Property and Selection Mechanism

We took 72 photos of different combinations of real fruits displayed on a dining table. For better performance of SAM, we did not choose to display the stimuli on a cleaner virtual background, like a whiteboard, because the image distributions of

[4] See chapter one, section 2 for a full discussion about SAM.

real photos are closer to the image distributions of the training sample. Each image contains two piles of fruits, and each pile contains three to five fruits. We flipped all the images in the horizontal direction so that we got another 72 images with the same content, but with the pile locations flipped.[5] These 174 images consist of our image pool.

We selected 20 images from the image pool, and all of the selected images satisfy four conditions:

1. **One-side saliency centered**: All of the selected images are strictly one-side saliency centered, which means that the most salient locations only appear in one fruit pile. Figure 4.4.1c represents an example of a one-side salient image, while Figure 4.4.1a shows an image that is not one-side saliency centered. Formally, consider two sets of pixels constituting the left pile and the right pile, P_l and P_r. Function s, the saliency model, maps the union of P_l and P_r to $[0, 1]$. The most salient location of an image consists of a set of pixels $S_h : \{x|s(x) > 0.90\}$.[6] An image is one-side saliency centered, if and only if exactly one of the two conditions hold true: $S_h \cap P_l = \emptyset$ or $S_h \cap P_r = \emptyset$.

2. **Balanced saliency center locations**: All of the selected images have saliency centers equally located on the left side or right side. Half of the images have saliency centers on the left, and the other half have them on the right.

3. **Balanced valuation distribution**: There are only two types of fruits: orange and apple. Each apple is worth 1.3 dollars and each orange is worth 2.2 dollars.[7] The total value differences between the two piles range from 0.4 dollars to 4 dollars. There are exactly 50% of rounds with the more rewarding option located on the left and 50% of rounds with the more rewarding option on the right.

4. **Balanced congruences**: An image will be called "congruent" if the more rewarding option on this image is also the more salient option. Among all images, we have 50% congruent images and 50% incongruent ones. No image contains two piles with the same amount of values to avoid confusion.

[5] This procedure is to avoid any left-right biases when taking images. It is done using a Matlab function flipimg().

[6] 0.90 is a reasonable threshold we set for this particular experiment. It can be anything higher or lower depending on the stimuli.

[7] We did a pilot experiment with integer unit values. It turns out that integer values are too easy for the subjects where we do not see any variation in choices.

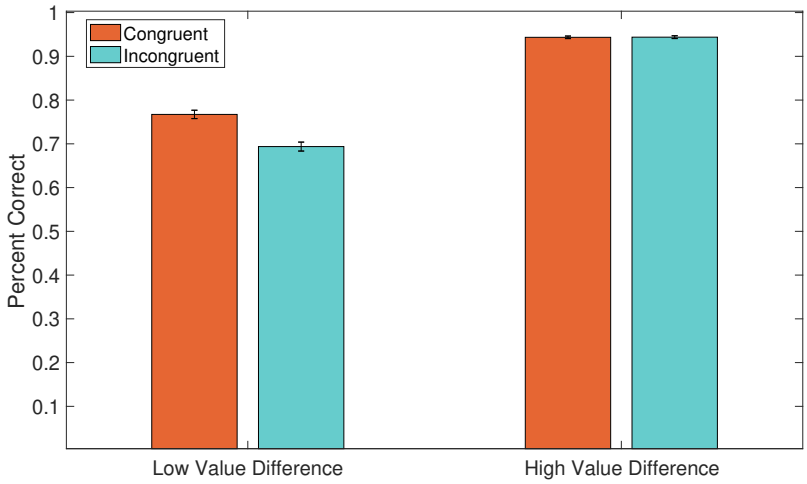

We split all stimuli into high valuation group and low valuation group with 10 images in each group. Correction rates are plotted accordingly between two conditions: congruent and incongruent within each group.

Figure 4.5.1: Effect on Correction Rate

5. **Balanced number of fruits**: Among the total 20 images, 18 images have the two fruit piles with only one unit in number differences. The last two images have two unit number differences. Eleven images have more fruits on the left, and nine images have more fruits on the right.

4.5 Results

More mistakes in incongruent conditions

When the most salient option is also the more rewarding option (the congruent condition), choice accuracy is 85%. When the salient option is not the highest-value option (incongruent condition), this level drops to 79%. This drop in accuracy is not large in magnitude, but is highly significant (p-value = 0.002). It is clear evidence that pure SAM-predicted salience has a causal effect on individual choice. The influence of salience is much larger, not surprisingly, when the value difference between the two choice sets is smaller rather than larger.[8]

Table 4.5.1 presents the regression results of choices and accuracy using all choice

[8]Furthermore, we split the total 20 images into high valuation difference set and low valuation difference set (cut-off = $0.55, each set contains 10 images) and compared the correction rate within each set (Figure 4.5.1). We found that the correction rate drops from 0.78 to 0.69 (p-value = 0.01) when the value differences between the two options are low. On the other hand, when value differences are high, people can make the right choice most of the time regardless of the saliency manipulation (correction rate = 0.94 for both conditions and p-value = 0.91).

data. Models (1) and (2) regress choices (y=1 indicates left pile) on value differences between two piles, whether a choice is salient (Saliency), and controls. The result indicates that valuation and saliency both contribute to the final choices. Subjects of course choose more rewarding options to win more money; they also have tendencies to choose more salient options and make mistakes. Models (3) and (4) regress the correctness of a choice on the absolute value difference between two alternatives, congruency property, and controls. The significant coefficient of congruency in regression (3) confirms the conditional treatment effect that people make more mistakes when the saliency feature conflicts with the reward feature conditioning on the same value difference. Furthermore, there is a weak negative interaction effect over value difference and congruency implying that saliency effect becomes weak when value difference is large. This result remains robust after clustering the standard deviation by individuals (see Appendix, Table 4..1).

One boundary condition is endogenous time allocation: When there is no time limit (N=25), participants in both conditions are near the ceiling of perfect accuracy (congruent =0.94 and incongruent=0.96).

Response time: Decisions are slower when saliency conflicts with reward

Other than choices, saliency and valuation might affect decisions on other domains like processing time. In this section, I will show the results of how saliency and valuation together affect decision times.

We did not find an unconditional treatment effect on response time between congruent trials and incongruent trials (p-value = 0.62). Subjects on average only spend 3.97s to finish the task ($\sigma = 2.7s$). Table 4.5.3 presents two regression models and the result implies a conditional treatment effect of congruency on response time. Model (1) regresses response time, (2) regresses normalized response time[9], both on absolute value difference between two piles (abs(ValueDiff)), absolute fruit number difference (abs(ItemsMore)), whether congruent (Congruency), interaction between value difference and congruency, and controls. Both models found a significant negative effect on response time by varying the congruency condition. In other words, when the more salient item is also the more rewarding item, subjects can make a faster decision. Such an effect becomes weaker when the value difference becomes larger (positive interaction coefficient and negative congruency coefficient).

[9]Normalized on each person's average response time

	Dependent variable:			
	Choice		Correctness	
	(1)	(2)	(3)	(4)
Value Difference		1.575***		
		(0.112)		
saliency	0.263**	0.485***		
	(0.111)	(0.143)		
congruency			0.472***	0.790***
			(0.147)	(0.212)
abs(valueDiff)				0.947***
				(0.155)
congruency:abs(valueDiff)				−0.365*
				(0.214)
gender:Male	0.066	0.120	−0.114	−0.124
	(0.116)	(0.148)	(0.152)	(0.158)
education:Bachelor	−0.451*	−0.714**	−0.131	−0.093
	(0.266)	(0.350)	(0.372)	(0.384)
education:High School	−0.425	−0.685**	−0.347	−0.329
	(0.262)	(0.345)	(0.365)	(0.376)
education:Master	−0.282	−0.457	−0.048	0.0002
	(0.297)	(0.388)	(0.415)	(0.428)
income:10k to 50k	−0.026	0.027	−0.046	−0.064
	(0.163)	(0.206)	(0.207)	(0.214)
income:50k to 100k	0.067	0.086	0.040	−0.006
	(0.212)	(0.271)	(0.277)	(0.287)
income:200k to 200k	0.262	0.517*	0.839**	0.869**
	(0.244)	(0.312)	(0.384)	(0.393)
income:200k and above	0.165	0.469	−0.354	−0.321
	(0.294)	(0.369)	(0.353)	(0.367)
Constant	0.314	0.180	1.557***	0.662
	(0.307)	(0.398)	(0.418)	(0.443)
Observations	1,309	1,309	1,307	1,307
Log Likelihood	-898.679	-585.019	-607.871	-556.045
Akaike Inf. Crit.	1,817.359	1,192.038	1,235.741	1,136.090
Note:				*p<0.1; **p<0.05; ***p<0.01

Saliency: whether an option is more salient. Value Difference: left value total - right value total. abs(ValueDiff) is the absolute term of value differences.

Table 4.5.1: Regression for Saliency on Choice.

	Dependent variable:	
	Response Time	Normalized Response Time
	(1)	(2)
abs(ValueDiff)	−0.725***	−0.176***
	(0.066)	(0.014)
Congruency	−0.575***	−0.123***
	(0.194)	(0.040)
abs(valueDiff):congruency	0.342***	0.079***
	(0.096)	(0.020)
gender:Male	−0.019	0.001
	(0.137)	(0.028)
income:10k to 50k	−0.285	0.005
	(0.193)	(0.039)
income:50k to 100k	−0.215	0.014
	(0.251)	(0.051)
income:200k to 200k	0.063	0.003
	(0.286)	(0.059)
income:200k and above	−0.939***	−0.011
	(0.349)	(0.071)
education:Bachelor	0.313	−0.016
	(0.307)	(0.063)
education:High School	−0.300	−0.013
	(0.303)	(0.062)
education:Master	−0.110	−0.017
	(0.344)	(0.070)
Constant	6.045***	1.365***
	(0.443)	(0.090)
Observations	1,309	1,309
R^2	0.119	0.144
Adjusted R^2	0.111	0.136
Residual Std. Error (df = 1296)	2.384	0.487
F Statistic (df = 12; 1296)	14.641***	18.145***
Note:		*p<0.1; **p<0.05; ***p<0.01

All variables are the same as table 4.5.1 except that y becomes response time and normalized response time normalized on a person level.

Table 4.5.3: Regression for Saliency on Response Time.

4.6 Results in a more complicated environment

We further test this paradigm in a more complex design (Figure 4.6.1), under which it will be almost impossible to actually calculate the valuation even if you put enough attention effort.

New Design

N=25 subjects on prolific participated in the new paradigm session and none of them has participated in any of the previous sessions. Comparing to the previous design where each alternative contained only 3 to 5 fruits, now each alternative contains 8 to 15 fruits. And the total number of types rise to six types. We follow the same stimulus selection mechanism and payment schema as the one described in section 4.4. The flow of the experiment also resembles the previous one described in figure 4.3.1 except that now the subjects have longer time to decide for each trial (20s), and the unit value of the fruits will always be written on the screen in case they forget.

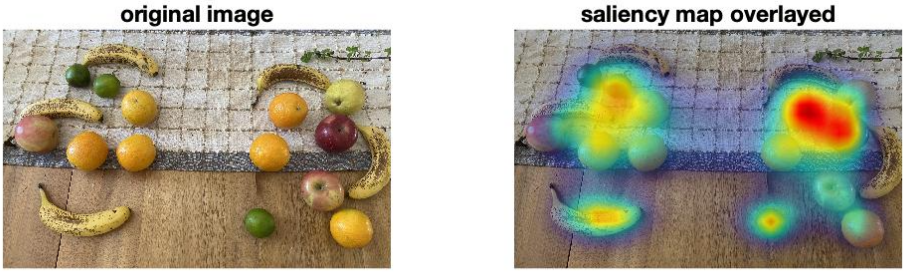

Figure 4.6.1: A more complex design.

The value of each fruits: apple,4 ;orange,3;banana,7;lime,5 ;pear,3; mango,6

Results

We also find an strong unconditional treatment effect on this new session between congruency and correctness (p-value = $4 * 10^{-5}$). The conditional treatment effect through a model similar to Table 2.2 is summarized in Table 4.6.1. It shows that more congruent choice environment, where the more rewarding item is also labeled salient, can significantly enhance the probability of choosing the more rewarding option even when the ground truth condition for rewards is more unclear under the current design.

Different from before, we found no effect on response time both for an unconditional treatment effect and a conditional treatment effect (regression analysis). Subjects

now spend on average 7.9s (σ = 5.0s) on each trial with a 20s time limit. We can see that for both experiments, no matter whether it is hard or easy, subjects are far from using up their time limit even though they still make mistakes.

	Dependent variable: correctness			
	(1)	(2)	(3)	(4)
Congruency	0.979***	0.867***		0.985***
	(0.343)	(0.213)		(0.219)
ValueDiff	0.103***		0.082***	0.103***
	(0.036)		(0.029)	(0.030)
Interaction	0.001			
	(0.066)			
Constant	−0.071	0.457***	0.383**	−0.073
	(0.214)	(0.120)	(0.162)	(0.192)
Observations	489	489	489	489
Log Likelihood	-290.304	-296.601	-301.186	-290.304
Akaike Inf. Crit.	588.607	597.203	606.371	586.608
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 4.6.1: Regression models on correction rate

All the results remain robust after clustering the standard deviations on an individual basis.

4.7 Conclusion and discussion

This chapter demonstrates that bottom-up visual saliency remains important for reward-based decision making, even if the reward criterion was predetermined and well-acknowledged. People keep making mistakes even when a more rewarding opportunity is not hard at all to distinguish. We find that saliency is a key feature leading to such mistakes that people make more mistakes when saliency is in the opposite direction of the more rewarding option. An algorithmic approach, saliency attentive map (SAM), remains to be powerful that it can enhance the correction rate by increasing the saliency level measured by SAM. Such an effect remains robust when the choice environment becomes more complicated. But it will go away when people have unlimited time to consider.

All the experiment findings fit well with a novel rational inattention model, the optimal solution of which predicts that saliency enhances corrected choice when the saliency property of choice is consistent with its reward property. Such bias only appears when the marginal benefit of learning is small.

Discussions

This model sheds light on how to understand saliency, reward, and mistakes in an incentivized mechanism. The result is well aligned with the selection theory, which is trending now in cognitive science (Failing and Theeuwes, 2018; Awh, Belopolsky, and Theeuwes, 2012) doubting the traditional dichotomy approach by separating "bottom-up" and "top-down"(Itti and Koch, 2001; Frintrop, Rome, and Christensen, 2010; Veale, Hafed, and Yoshida, 2017). The selection theory argues that the visual attention mechanism originates from a reward-based history effect under a biological view. The selection theory also reflects in the short-term. Associations between stimuli and rewards imply that memory processes are capable of shaping the attentional selection. Priming is an example that the current visual attention will be biased by stimuli that were selected in a recent past (Della Libera and Chelazzi, 2009). A complete model should include the entire bilateral relationship between the reward/goal and the saliency.

This study would be more complete if we applied eye-tracking techniques. Due to the pandemic of COVID-19[10], we were not able to do any in-lab experiment. This task by itself will be a suitable paradigm to test the task relevancy top-down control and compare the result to the bottom-up driven saliency. Task relevance has been identified as the primary driver of attention in many of the natural tasks before (Hayhoe, 2000; Land, Mennie, and Rusted, 1999; Sprague, Ballard, and Robinson, 2007). It is also likely to be the case in our fruit experiment. What we did in this chapter is just to evaluate the effect coming from the feature-driven saliency. For instance, the true attention patterns might migrate from the bottom-up mode to the search mode that, in each trial, subjects are "top-down" searching for the most valuable item. In our case, the more valuable item is an orange, which is not more visually salient comparing to an apple. Therefore, measuring the fixation change on an object basis will be a critical move for future studies. Such a change of attention patterns may lead to the dynamic change of decision strategies accordingly.

References

Awh, Edward, Artem V. Belopolsky, and Jan Theeuwes (Aug. 2012). "Top-down versus bottom-up attentional control: a failed theoretical dichotomy". In: *Trends in Cognitive Sciences* 16(8), pp. 437–443. ISSN: 1364-6613. DOI: 10.1016/

[10]At the time of this dissertation, COVID-19 outbreaks and all in lab works are terminated, including eye-tracking experiments.

j.tics.2012.06.010. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3426354/ (visited on 05/30/2018).

Caplin, Andrew and Mark Dean (2013). *Behavioral implications of rational inattention with shannon entropy*. Tech. rep. National Bureau of Economic Research.

Caplin, Andrew and Mark Dean (2015). "Revealed preference, rational inattention, and costly information acquisition". In: *American Economic Review* 105(7), pp. 2183–2203.

Caplin, Andrew, Mark Dean, and John Leahy (2019). "Rational inattention, optimal consideration sets, and stochastic choice". In: *The Review of Economic Studies* 86(3), pp. 1061–1094.

Cornia, Marcella et al. (2016). "Predicting human eye fixations via an LSTM-based saliency attentive model". In: *arXiv preprint arXiv:1611.09571*.

Della Libera, Chiara and Leonardo Chelazzi (2009). "Learning to attend and to ignore is a matter of gains and losses". In: *Psychological science* 20(6), pp. 778–784.

Failing, Michel and Jan Theeuwes (2018). "Selection history: How reward modulates selectivity of visual attention". In: *Psychonomic bulletin & review* 25(2), pp. 514–538.

Frintrop, Simone, Erich Rome, and Henrik I. Christensen (2010). "Computational visual attention systems and their cognitive foundations: A survey". In: *ACM Transactions on Applied Perception (TAP)* 7(1), p. 6.

Hayhoe, Mary (2000). "Vision using routines: A functional account of vision". In: *Visual Cognition* 7(1-3), pp. 43–64.

Itti, Laurent and Christof Koch (2001). "Computational modelling of visual attention". In: *Nature Reviews Neuroscience* 2(3), p. 194.

Kamenica, Emir and Matthew Gentzkow (2011). "Bayesian persuasion". In: *American Economic Review* 101(6), pp. 2590–2615.

Krajbich, Ian, Carrie Armel, and Antonio Rangel (2010). "Visual fixations and the computation and comparison of value in simple choice". In: *Nature Neuroscience* 13(10), p. 1292.

Land, Michael, Neil Mennie, and Jennifer Rusted (1999). "The roles of vision and eye movements in the control of activities of daily living". In: *Perception* 28(11), pp. 1311–1328.

Milosavljevic, Milica et al. (2012). "Relative visual saliency differences induce sizable bias in consumer choice". In: *Journal of Consumer Psychology* 22(1), pp. 67–74.

Orquin, Jacob L and Simone Mueller Loose (2013). "Attention and choice: A review on eye movements in decision making". In: *Acta psychologica* 144(1), pp. 190–206.

Scarpa, Riccardo et al. (2013). “Inferred and stated attribute non-attendance in food choice experiments”. In: *American Journal of Agricultural Economics* 95(1), pp. 165–180.

Shimojo, Shinsuke et al. (2003). “Gaze bias both reflects and influences preference”. In: *Nature neuroscience* 6(12), pp. 1317–1322.

Sims, Christopher A (2003). “Implications of rational inattention”. In: *Journal of monetary Economics* 50(3), pp. 665–690.

Sims, Christopher A (2006). “Rational inattention: Beyond the linear-quadratic case”. In: *American Economic Review* 96(2), pp. 158–163.

Sprague, Nathan, Dana Ballard, and Al Robinson (2007). “Modeling embodied visual behaviors”. In: *ACM Transactions on Applied Perception (TAP)* 4(2), 11–es.

Towal, R Blythe, Milica Mormann, and Christof Koch (2013). “Simultaneous modeling of visual saliency and value computation improves predictions of economic choice”. In: *Proceedings of the National Academy of Sciences* 110(40), E3858–E3867.

Veale, Richard, Ziad M Hafed, and Masatoshi Yoshida (2017). “How is visual salience computed in the brain? Insights from behaviour, neurobiology and modelling”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 372(1714), p. 20160113.

APPENDIX

	Dependent variable: Correctness	
	(1)	(2)
Congruency	0.472***	0.539***
	(0.156)	(0.156)
abs(ValueDiff)		0.793***
		(0.057)
gender:Male	−0.114	−0.122
	(0.321)	(0.321)
income:10k to 50k	−0.046	−0.063
	(0.439)	(0.439)
income:50k to 100k	0.040	−0.005
	(0.564)	(0.564)
income:200k to 200k	0.839	0.868
	(0.592)	(0.592)
income:200k and above	−0.354	−0.326
	(1.115)	(1.115)
education:Bachelor	−0.131	−0.094
	(0.873)	(0.873)
education:High School	−0.347	−0.328
	(0.855)	(0.855)
education:Master	−0.048	0.001
	(0.960)	(0.960)
Constant	1.557	0.763
	(1.086)	(1.086)
Note:	*p<0.1; **p<0.05; ***p<0.01	

This model is the robust version of table 4.5.1

Table 4..1: The robust regression on correction rate.

Chapter 5

DISCUSSION AND FUTURE DIRECTIONS

5.1 Conclusion

In summary, all dissertation chapters are trying to draw connections among the computational model of visual saliency, the eye fixation transitions, and the final choice outcome and make a congruent and complete attention-choice mechanism based on it.

Chapter four focuses on the role of attention bias in a binary choice problem. It firstly developed a saliency-enhanced rational inattention model by incorporating attention factors in the traditional RI model. This theory provides a new perspective to understand our visual system and the biological meaning of the saliency map. In this new model, part of the information on reward outcome is given for free (no information cost). The solution of this new model leads to a common phenomenon that people tend to think a more salient option is also a more rewarding option. To test this idea using experiments, we successfully separated the reward property and the attention property of choice alternatives. We let subjects select between two choices and found that the visual saliency property can enhance the choice correction rates when the more rewarding alternative is also labeled salient.

The other two chapters answered another vital question that is how visual saliency affects strategic decision makings. People have realized the existence of the "focal point" in games for long but lack a scientific model frame to build up a testable model. We build up a testable model by borrowing the computational model of visual saliency (SAM). Applying the SAM algorithm solves three problems. First, it predicts with substantial accuracy how often people match on different images. Second, when predicted salience is embedded in a cognitive hierarchy or level-k model, the model generates a seeker's advantage. Salience strength estimated from hider-seeker games can also "portably" predict choices in matching games that have a different structure. Third, using eye-tracking data and salience, an HMM predicts numerically what is likely to happen under time pressure. The predicted seeker win rate under 2s pressure is 14%, which is not too far from the observed win rate with feedback (15%) and is higher than the 6s win rate of 9%.

5.2 Limitation

One major limitation of this dissertation, like said before, is that we narrow the big concept of attention to be only the visual perception biases. The entire attention concept in decision making definitely goes beyond the visual world. Using the eye fixation map as a metric for attention is, of course, a more direct measurement than to infer it from the choice data. More direct measurement will be a direct set of weights on all choice alternatives when people are doing the value computations. Such a measurement will be close enough to the "attention" concept in any decision theory. Unfortunately, our current technology is very far from being able to execute this perfect measurement.

Even within the visual saliency range, what we know now in the visual system and the brain is very limited. We already have good computational models for predicting early gaze patterns using visual scenes. This is usually considered as the "bottom-up" approach in vision neuroscience though the current deep neural network method has already included many top-down features like meaningful semantic segments. Other "top-down" methods are more separated from the line of the image-driven methods. For example, Henderson and Hayes (2017) developed a meaning map method which uses small image patches manually labeled by human labelers. Other studies on the "top-down" perspective go under the topic of visual search, which is another powerful mechanism of gaze redirection (Itti and Baldi, 2009; Itti and Koch, 2000; Navalpakkam and Itti, 2007). It is controlled by higher-order brain areas, including the frontal lobes. These brain areas connect back into the visual cortex and early visual areas to form an entire saliency network (Itti and Koch, 2001). There is no congruent model yet that can combine these two pathways. Therefore, a significant part of the "top-down" guided attention is not under control using our method.

Another important question that goes beyond the scope of this dissertation is whether there exists a feedback loop where the attention mechanism is also affected by the reward value distribution. From an evolutionary point of view, the saliency mechanisms evolve towards the historically rewarding features (Failing and Theeuwes, 2018). This part of the impact has already been included in the feature coded algorithm. In short-term, associations between stimuli and rewards imply that memory processes are capable of shaping the attentional selection. Priming is an example that current visual attention will be biased by stimuli that were selected in a recent past (Della Libera and Chelazzi, 2009). A complete model should include

the entire bilateral relationship between the reward/goal and the saliency.

5.3 Future directions

One possible future direction for simple choice decision making problems is to study time-variant models of the external attention attributes and choices. For example, can people resist the bias of attention by actively slowing down the decision process? Though the aDDM model provides quantitative predictions on response time distribution. It is on a passive point of view that decision time is determined intrinsically by the decision mechanism. An extended question on this would be whether people can consciously notice the potential loss of this bias and actively control it. Chapter 4 provides a possible way to do it. One can continue the paradigm of Chapter 4 and calculate the utility equivalence for the attentional biases.

Another continuation of this work could be along the direction of rational inattention models and see whether there is any biological confound in the brain corresponding to this model. The rational inattention model is the model for mistakes. Attention biases can also cause errors. A self-monitoring feedback loop has been recognized during the single-unit recording technique in the dorsal anterior cingulate cortex (dACC) area (Fu et al., 2019). A possible biological basis for the attention network might also exist in this area. If the people self-monitor such attention biases, a bilateral model between attention and valuation would be closer to the truth.

The hidden Markov model (HMM) we proposed for modeling the choice process data also has much potential to develop. It can be as general as a drift diffusion model. Markov models are well connected to the drift diffusion model class if we view each position during DDM as a hidden state. The HMM model has many advantages. The hidden structure in HMM from the hidden states to the observations will make the model flexible to be applied on any given number of choice problems. One crucial question remaining will be how to choose the best number of hidden states.

For the strategic part, we only provide an experimental work of the attention model. A related use case will be worthwhile to explore in the future. For example, strategic competition often happens in information disclosure activity. Sometimes, information disclosure can be viewed as a matching game. For example, both the shopping mall and the shoppers want to know that there is a discount happening. But sometimes it becomes a hider-seeker game that the shopping mall does not want the user to notice that the promotions only apply to a small group of products.

Therefore, we often see that the information about the former case is written in the focal place while the information of the latter case is often hidden in obscured places. A testable data set will be useful for the application of our model.

References

Della Libera, Chiara and Leonardo Chelazzi (2009). "Learning to attend and to ignore is a matter of gains and losses". In: *Psychological science* 20(6), pp. 778–784.

Failing, Michel and Jan Theeuwes (2018). "Selection history: How reward modulates selectivity of visual attention". In: *Psychonomic bulletin & review* 25(2), pp. 514–538.

Fu, Zhongzheng et al. (2019). "Single-neuron correlates of error monitoring and post-error adjustments in human medial frontal cortex". In: *Neuron* 101(1), pp. 165–177.

Henderson, John M. and Taylor R. Hayes (2017). "Meaning-based guidance of attention in scenes as revealed by meaning maps". In: *Nature Human Behaviour* 1(10), p. 743.

Itti, Laurent and Pierre Baldi (2009). "Bayesian surprise attracts human attention". In: *Vision Research* 49(10), pp. 1295–1306.

Itti, Laurent and Christof Koch (2000). "A saliency-based search mechanism for overt and covert shifts of visual attention". In: *Vision research* 40(10-12), pp. 1489–1506.

Itti, Laurent and Christof Koch (2001). "Computational modelling of visual attention". In: *Nature Reviews Neuroscience* 2(3), p. 194.

Navalpakkam, Vidhya and Laurent Itti (2007). "Search goal tunes visual features optimally". In: *Neuron* 53(4), pp. 605–617.

BIBLIOGRAPHY

Baum, Leonard E et al. (1970). “A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains”. In: *The Annals of Mathematical Statistics* 41(1), pp. 164–171.

Crawford, Vincent P. and Nagore Iriberri (2007a). “Fatal attraction: Salience, naivete, and sophistication in experimental" hide-and-seek" games”. In: *American Economic Review* 97(5), pp. 1731–1750.

Crawford, Vincent P. and Nagore Iriberri (2007b). “Level–k Auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private–value auctions?” In: *Econometrica* 75(6), pp. 1721–1770.

Murphy, Kevin (2005). *Hidden Markov Model (HMM) toolbox for Matlab Retrieved May 25 2012*.

Nagel, Rosemarie (1995). “Unraveling in guessing games: An experimental study”. In: *The American Economic Review* 85(5), pp. 1313–1326.