Rational Inattention, Optimal Consideration Sets and Stochastic Choice*

Andrew Caplin[†], Mark Dean[‡], and John Leahy[§]

June 2016

Abstract

We unite two of the most common approaches to modelling limited attention in choice by showing that the rational inattention model (Sims [2003]) implies the formation of consideration sets – only a subset of the available alternatives will be considered for choice. We provide a new set of necessary and sufficient conditions for rationally inattentive behavior which allow the identification of consideration sets. In simple settings, consideration sets are based on a cutoff strategy on expected ex–ante utility modified by attention costs. The makeup of consideration sets can change in highly non-monotone ways with attention costs and risk aversion.

1 Introduction

Attention is a scarce resource. The impact of attentional limits has been identified in many important economic settings,[1] leading to widespread efforts to model the effect of such con-

*We thank Dirk Bergemann, Henrique de Oliveira, Xavier Gabaix, Sen Geng, Andrei Gomberg, Daniel Martin, Filip Matejka, Alisdair McKay, Stephen Morris, and Michael Woodford for their constructive contributions. This paper develops some of the concepts of Caplin and Dean [2013] and subsumes those parts of that paper that are common.

†Center for Experimental Social Science and Department of Economics, New York University. Email: andrew.caplin@nyu.edu

‡Department of Economics, Columbia University. Email: mark.dean@columbia.edu

§Department of Economics, University of Michigan and NBER. Email: jvleahy@umich.edu

[1]For example, shoppers may buy unnecessarily expensive products due to their failure to notice whether or not sales tax is included in stated prices (Chetty *et al.* [2009]). Buyers of second-hand cars focus their attention on the leftmost digit of the odometer (Lacetera *et al.* [2012]). Purchasers limit their attention to a relatively small number of websites when buying over the internet (Santos *et al.* [2012]).

straints.

One key implication of limited attention is that a decision maker may consider only a subset of the available alternatives, ignoring all others. The concept of 'consideration sets' has a long history in the marketing literature, which extensively demonstrates their importance.[2] More recently, economists have begun to understand the importance of consideration sets for many areas of study - including revealed preference analysis, the price setting behavior of firms, and demand estimation.[3]

A second implication of attentional constraints is that choice may be stochastic: a decision maker may make different choices in seemingly identical situations. Random choice has been demonstrated in a wide variety of experimental settings.[4] The relationship between stochasticity and attention constrained choice has been emphasized by the 'rational inattention' approach of Sims [2003], in which the decision maker (DM) chooses information optimally given their decision problem, with costs based on the Shannon mutual information between prior and posterior beliefs (henceforth the Shannon model).

In this paper we unite these two approaches, and present a theory of optimal consideration set formation based on rational inattention. We show that an implication of the Shannon model in discrete choice settings is that, typically, many options will receive no consideration and will never be chosen.[5] The set of considered alternatives arises endogenously, based on prior beliefs and attention costs. Moreover, the same parameters determine a pattern of stochastic choice 'mistakes' amongst considered alternatives, in line with experimental findings.[6] Unlike current models, our approach therefore provides a tractable, parsimonious model of both endogenous consideration set formation and choice mistakes within the consideration set.

In order to develop our model, we introduce a novel set of necessary and sufficient first order conditions for the Shannon model. These build on the necessary conditions of Matejka and McKay [2015] (henceforth MM), who characterize the pattern of stochastic choice implied by the Shannon model *amongst alternatives which are chosen with positive probability.* We introduce a set of easy-to-check inequality constraints which determine the set of chosen actions, and so also the set of actions which are never chosen. These conditions are crucial to the solution of the Shannon model in any application - not only those we consider in

[2] For example Hoyer [1984], Hauser and Wernerfelt [1990] and Roberts and Lattin [1991].

[3] See for example Ching *et al.* [2009], Eliaz and Spiegler [2011], Caplin *et al.* [2011], Masatlioglu *et al.* [2012], De Clippel *et al.* [2014], Manzini and Mariotti [2014].

[4] See for example Mosteller and Nogee [1951], Reutskaja *et al.* [2011] and Agranov and Ortoleva [2015].

[5] See Matejka and Sims [2011].

[6] See for example Geng [2016]

this paper: generally not all actions with be taken at the optimum, and there will be many suboptimal patterns of behavior which satisfy MM's conditions amongst different potential consideration sets.

Using our necessary and sufficient conditions, we consider behavior in three variants of the standard consumer problem. In each case, the consumer must choose one of a set of available alternatives. The value of each alternative is ex ante uncertain, but that uncertainty can be reduced by allocating attention and paying the associated costs. The three decision making environments vary in the assumed correlation structure between the valuations of different alternatives.

In our first application the consumer must choose between a number of alternatives, one of which is of high quality and the remainder of which are of low quality. The identity of the high quality alternative is unknown ex ante, but can be learned by the consumer. In this setting, the Shannon model implies that consideration sets are determined by a threshold strategy: consumers will pay attention only to alternatives which have a prior probability of being high quality which is above an endogenously determined threshold. Alternatives below this threshold will never be chosen even though there is a chance that this set includes the high quality good. Amongst considered alternatives, attention is allocated in such a way that ex post all choices are identical: the probability of any alternative being of high quality conditional on being chosen is the same, regardless of prior belief. This feature highlights an important link between the Shannon model and the Drift Diffusion Model (DDM) of information acquisition and choice, which is heavily studied in psychology.

In our second environment we assume that the valuation of different alternatives are distributed independently. In this setting, consideration set formation is again driven by a cutoff strategy. However, the ranking of alternatives is now determined by the expectation of a convex transformation of the payoffs. This transformation reflects the gains to information acquisition: the ability to take an action when its payoff is relatively high and avoid it when its payoff is low. Given this convex transformation, the composition of consideration sets can change in surprising and non-monotonic ways as information costs change. We provide an example in which the consumer must choose either a safe alternative, the value of which is known ex ante, or one of a set of risky alternatives, the value of which must be learned. The safe alternative only appears in the consideration at very high information costs - when it allows the consumer to be uninformed, or very low information costs - when it is chosen if all the risky alternatives turn out to be of low quality. For information costs in an intermediate range, the consideration set consists only of the risky alternatives. We further demonstrate that, if prizes are denominated in monetary terms, the make up of the consideration set is jointly determined by the information costs and risk aversion of the consumer: risky alternatives will only be used if the consumer has low attention costs or low risk aversion.

Finally, we look at the most general case of arbitrary correlation between the valuations of different alternatives - for example of the type that might occur in the choice of financial products. In this case, no simple cutoff rule determines the consideration set. This is due to the fact that even risk neutral consumers have a hedging motive in this environment: the value of act depends on its payoff relative to other acts. Given this hedging motive, the consideration set will depend on the complete range of available alternatives. We show that our conditions imply a simple test of whether a new alternative will be considered if it is introduced to an existing market, and identify the lowest cost way of ensuring such an alternative will be chosen.

Section 4 discusses the relationship of our work to existing models of consideration set formation. Recent papers have typically taken consideration sets as primitive for the consumer (in the same way that preferences are primitive), therefore sidestepping the issue of how the set of considered items is determined. These papers then focus on identifying the consideration set from choice behavior (Masatlioglu *et al.* [2012], Manzini and Mariotti [2014]) or on understanding firm behavior conditional on such sets (Eliaz and Spiegler [2011]). Moreover, these models assume that decision makers deterministically maximize preferences on the consideration set. Yet recent evidence (e.g. Geng [2016]) shows that choice may be stochastic even amongst considered alternatives. As with Manzini and Mariotti [2014], our work provides a link between the study of consideration sets and the recent literature aimed at understanding stochastic choice data (e.g. Agranov and Ortoleva [2015], Manzini and Mariotti [2016], Apesteguia and Ballester [2016]). An earlier literature in marketing discussed models of endogenous consideration set formation (e.g. Hauser and Wernerfelt [1990], Roberts and Lattin [1991]). However, these have typically had to focus on very specific cases for the sake of tractability: sequential choice of information in more complex settings quickly become intractable (Gabaix *et al.* [2006]).

Section 2 introduces the Shannon model and our necessary and sufficient conditions. Section 3 describes our three applications to consumer search. Section 4 reviews the existing literature, and section 5 concludes.

2 The Shannon Model

We consider a consumer who faces a decision problem which consists of a number of different alternatives from which they must make a choice. The value of each alternative is determined by the underlying state of the world. Prior to choice, the decision maker can receive information about the state of the world in the form of an *information structure,* which consists of a set of signals, and a stochastic mapping between the true state of the world and these signals. More accurate signals will lead to better choices, but are more costly, with costs based on the Shannon mutual information between prior and posterior beliefs. This is the model of 'rational inattention' introduced by Sims [2003].

2.1 The Decision Making Environment

There is a finite set of possible states of the world Ω, with $\omega \in \Omega$ denoting a generic state. An action is a mapping from states of the world to utilities. We use $\mathcal{A}$ to denote the set of actions, and $u : \mathcal{A} \times \Omega \to \mathbb{R}$ to identify the utility of each action in each state. A decision problem (μ, A) consists of a prior distribution $\mu \in \Delta(\Omega)$ over these states of the world and a finite subset of options $A \subseteq \mathcal{A}$ from which the decision maker must choose.

It is well known that, for the Shannon model one can directly model the probability of choosing each action in each state, rather than the choice of information structure (see for example MM). Thus, given the decision problem (μ, A), the decision maker chooses the probability of receiving option $a \in A$ in each state ω: i.e., for a given A they choose a $P : \Omega \to \Delta(A)$, with $P(a|\omega)$ denoting the probability of choosing action a in state ω, and $\mathcal{P}$ the set of all such state dependent stochastic choice functions.

The value of $P \in \mathcal{P}$ is given by the expected value of the actions chosen, minus information costs. These costs are based on the Shannon mutual information between states and actions - i.e. the difference between the expected entropy of the conditional and unconditional choice distributions:[7] Intuitively, having choice distributions which vary a lot with the state requires costly information. MM show that the optimization problem of a consumer

[7] The entropy of a distribution p on Ω is given by $\sum_{\omega \in \Omega} p(\omega) \ln p(\omega)$. Broadly speaking a high degree of entropy means that there is a lot of uncertainty.

facing a decision problem (μ, A) can be written as follows:

$$\begin{aligned} V(\mu, A) &= \max_{P \in \mathcal{P}} \sum_{\omega \in \Omega} \mu(\omega) \left(\sum_{a \in A} P(a|\omega) u(a, \omega) \right) \\ &\quad -\lambda \left[\sum_{\omega \in \Omega} \mu(\omega) \left(\sum_{a \in A} P(a|\omega) \ln P(a|\omega) \right) - \sum_{a \in A} P(a) \ln P(a) \right] \end{aligned} \tag{1}$$

where $P(a) = \sum_{\omega \in \Omega} \mu(\omega) P(a|\omega)$.

The first term is the expected payoff due to the set of state contingent choice probabilities. The second term is the cost of information which is equal to the mutual information between states and actions multiplied by the parameter λ, which describes the marginal cost of information.

2.2 Necessary and Sufficient Conditions

MM show that the optimal policy must satisfy the following condition for all actions $a \in A$ such that $P(a) > 0$:

$$P(a|\omega) = \frac{P(a) z(a, \omega)}{\sum_{b \in A} P(b) z(b, \omega)} \tag{2}$$

where $z(a, \omega) \equiv \exp(u(a, \omega)/\lambda)$. This condition states that the optimal policy "twists" the state-contingent choice probabilities towards states in which the payoffs are high.

While the conditions (2) are necessary, they are not sufficient. For example, for any option a, the stochastic choice function $P(a|\omega) = 1$ for all ω satisfies these conditions. This is not surprising since this is the optimal policy when only option a is available and the necessary conditions depend on the choice set only in the sense that all of the choices in the sum in the denominator must be available.

The key problem is that, while these necessary conditions determine stochastic choice amongst alternatives which are chosen with positive probability, they do not identify which alternatives belong to this set. Doing so is important because, as we shall see, generally there will be many unchosen actions in a given decision problem. We will describe the set of actions which are chosen with positive probability as the *consideration set,* and for every $P \in \mathcal{P}$ we will use $B(P)$ to denote the associated consideration set - i.e.

$$B(P) = \{a \in A | P(a) > 0\}$$

According to (2) it is sufficient to solve for the unconditional probabilities of the options, $P(a)$, as the state contingent choice probabilities, $P(a|\omega)$, are completely determined by these unconditional probabilities. To obtain conditions which are both necessary and sufficient we focus on the choice of these $P(a)$s. We can reformulate the maximization problem in terms of these unconditional probabilities. Substituting for $P(a|\omega)$ in the objective function described in equation (1) gives the following expression[8]

$$\sum_{\omega\in\Omega}\lambda\ln\left(\sum_{a\in A}P(a)z(a,\omega)\right)\mu(\omega),$$

meaning we can write a Lagrangian for problem:

$$\max_{P(a)}\sum_{\omega\in\Omega}\lambda\ln\left(\sum_{a\in A}P(a)z(a,\omega)\right)\mu(\omega)-\varphi\left(\sum_{a\in A}P(a)-1\right)+\sum_{a\in A}\xi^a P(a) \tag{3}$$

where φ is the Lagrangian multiplier on the constraint that the unconditional probabilities $P(a)$ must sum to 1, and ξ^a is the multiplier on the non-negativity constraint for $P(a)$.

The associated first order condition is

$$\lambda\sum_{\omega\in\Omega}\frac{z(a,\omega)}{\left(\sum_{b\in A}P(b)z(b,\omega)\right)}\mu(\omega)-\varphi+\xi^a=0$$

We can simplify further by noting that $\sum_{\omega\in\Omega}P(a|\omega)\mu(\omega)\equiv P(a)$ and (2) imply

$$P(a)=\sum_{\omega\in\Omega}P(a|\omega)\mu(\omega)=\sum_{\omega\in\Omega}\frac{P(a)z(a,\omega)}{\sum_{b\in A}P(b)z(b,\omega)}\mu(\omega)$$

or

$$\sum_{\omega\in\Omega}\frac{z(a,\omega)\mu(\omega)}{\sum_{b\in A}P(b)z(b,\omega)}=1,$$

whenever $P(a)>0$. It follows immediately that $\varphi=\lambda$. The first order condition becomes

$$\sum_{\omega\in\Omega}\frac{z(a,\omega)\mu(\omega)}{\sum_{b\in A}P(b)z(b,\omega)}\leq 1, \tag{4}$$

with equality if the option is chosen. These are the necessary and sufficient conditions for an optimum.

[8] See appendix for details.

Proposition 1 *The policy $P \in \mathcal{P}$ is optimal if and only if for all $a \in A$*

$$\sum_{\omega \in \Omega} \frac{z(a,\omega)\mu(\omega)}{\sum_{b \in A} P(b) z(b,\omega)} \leq 1$$

with equality if $a \in B(P)$, and for all such actions and states ω, $P(a|\omega)$ satisfies condition 2.

All proofs are relegated to the appendix.

We can gain some intuition for these conditions by considering the Blahut-Arimoto algorithm (see Cover and Thomas [2006]). The algorithm proceeds by first choosing $P(a|\omega)$ given a guess for the $P(a)$ and then generating a new $P(a)$ given $P(a|\omega)$. Since it can be shown that the objective function in equation (1) rises with each step, the algorithm converges. The solution to the first step is MM's necessary conditions

$$P_{n+1}(a|\omega) = \frac{P_n(a) z(a,\omega)}{\sum_{b \in A} P_n(b) z(b,\omega)}$$

The solution to the second step invokes Bayes rule: $P_{n+1}(a) = \sum_{\omega \in \Omega} P_{n+1}(a|\omega)\mu(\omega)$. Putting these together

$$P_{n+1}(a) = \left(\sum_{\omega \in \Omega} \frac{z(a,\omega)\mu(\omega)}{\sum_{b \in A} P_n(b) z(b,\omega)} \right) P_n(a)$$

The term in brackets is the left side of (4) and determines if $P(a)$ rises or falls. The algorithm can have one of two steady states. Either $P(a) > 0$ and the term in brackets is equal to one, a case that includes $P(a) = 1$, or the term in brackets is less than one and $P(a) = 0$ and can fall no further. (2) represents a twist in state dependent choice in the direction of the high payoff states. (4) ensures that these twists average out to one. If they don't then the probability of an action needs to be raised or lowered accordingly.

2.3 A Posterior-Based Approach

It is also informative to recast the solution of the Shannon model in terms of the implied posterior beliefs. Note any set of stochastic choices $P \in \mathcal{P}$ imply posterior belief $\gamma^a \in \Delta(\Omega)$ at any $a \in B(P)$. By Bayes rule,

$$\gamma^a(\omega) = \frac{P(a|\omega)\mu(\omega)}{P(a)},$$

where $\gamma^a(\omega)$ is the probability of state ω given the choice of a. There is a one-to-one mapping between the choice of state dependent stochastic choices $P \in \mathcal{P}$ and the choice of unconditional choice probabilities and posterior beliefs $\{\{P(a)\}_{a\in A}, \{\gamma^a\}_{a\in B(P)}\}$. We can rewrite the necessary and sufficient conditions of Proposition 1 in terms of these objects

It follows immediately from (2) that, for $a \in B(P)$

$$\frac{\gamma^a(\omega)}{z(a,\omega)} = \frac{\mu(\omega)}{\sum_{b\in A} P(b)z(b,\omega)}$$

The necessary and sufficient conditions (4) become

$$\frac{\gamma^a(\omega)}{z(a,\omega)} = \frac{\gamma^b(\omega)}{z(b,\omega)} \tag{5}$$

when options a and b are chosen and

$$\sum_{\omega\in\Omega} z(c,\omega)\frac{\gamma^a(\omega)}{z(a,\omega)} \leq 1 \tag{6}$$

when option a is chosen and option c is not. We therefore have the following alternative characterization of the optimal policy:

Proposition 2 *Consider the choice problem* (A, μ). *Consider a policy* $\{P^a\}_{a\in A}$ *and* $\{\gamma^a\}_{a\in B(P)}$. *The policy is optimal if and only if* $\sum_{a\in A} P(a)\gamma(\omega) = \mu(\omega)$ *and*

1. ***Invariant Likelihood Ratio (ILR) Equations for Chosen Options**: given* $a, b \in B(P)$, *and* $\omega \in \Omega$,
$$\frac{\gamma^a(\omega)}{z(a,\omega)} = \frac{\gamma^b(\omega)}{z(b,\omega)}$$

2. ***Likelihood Ratio Inequalities for Unchosen Options**: given* $a \in B(P)$ *and* $c \in A\backslash B(P)$,
$$\sum_{\omega\in\Omega} \left[\frac{\gamma^a(\omega)}{z(a,\omega)}\right] z(c,\omega) \leq 1. \tag{7}$$

3 Endogenous Consideration Set Formation

The necessary and sufficient conditions of Proposition 1 represent an important advance on the purely necessary condition of MM in terms of solving the Shannon model. They also highlight a particular feature of such solutions: the available actions can be divided into a *consideration set* of alternatives which are chosen with positive probability in every state of the world, and an excluded set of alternatives which are never chosen in any state. This provides a link between the Shannon model and models of choice with consideration sets which have long been popular in the marketing literature (see for example Hoyer [1984], Hauser and Wernerfelt [1990] and Roberts and Lattin [1991]).

In this section we explore this link further by applying the Shannon model to three variants of the consumer problem, and characterizing the resulting consideration sets in each case. We interpret the set A as representing the set of goods for sale, one of which the consumer will end up buying. The states of the world relate to the quality of each of the various goods.

The necessary and sufficient conditions are non-linear and so, in general, closed form solutions are not available. We consider three variants of the consumer problem that differ in their complexity. The first we can solve in closed form. In this case, the consumer knows that exactly one of the available goods is of high quality, while the rest are of low quality. While somewhat stylized as a consumer problem, this setting allows for a particularly clean understanding of the consideration set resulting from the Shannon model, and demonstrates an important link between the Shannon model and the Drift Diffusion Model which is popular in psychology. In the second case, we consider the 'classic' consumer problem in which the valuation of each alternative is drawn form an independent (but not necessarily identical) distribution. In this case we present a statistic that determines whether an action is or is not in the consideration set. Finally we discuss the case in which the valuation of different goods may be correlated.

3.1 Consumer Problem 1: Finding the Best Alternative

We begin with arguably the simplest possible case. The consumer is faced with a range of possible goods identified as a set $A = \{a_1, .., a_M\}$. One of these options is good, the others are bad. The utilities of the good and bad options are u_G and u_B, with $u_G > u_B$. The DM has a prior on which of the available options is good. We define the state space to be the same as the action space, $\Omega = A$, with the interpretation that state ω_i is the state in which option i is of high quality and all others are of low quality. Thus

$$\begin{aligned} u(a_i, \omega_j) &= u_G \text{ if } i = j \\ &= u_B \text{ otherwise.} \end{aligned}$$

$\mu(\omega_i)$ is therefore the prior probability that option a_i yields the good prize. Without loss of generality, we order states according to perceived likelihood with lower indexed states perceived as more likely,

$$\mu_i = \mu(\omega_i) \geq \mu(\omega_{i+1}) = \mu_{i+1}.$$

The DM can expend attentional effort to gain a closer understanding of where the prize is located. The cost of improved understanding is defined by the Shannon model with parameter $\lambda > 0$.

3.1.1 Characterization

To simplify characterization of the optimal strategy, it is convenient to transform parameters by defining x, $\delta > 0$ as below,

$$z(a_i, \omega_j) = \begin{cases} \exp\left(\frac{u_G}{\lambda}\right) \equiv x(1+\delta) & \text{if } i = j; \\ \exp\left(\frac{u_B}{\lambda}\right) \equiv x & \text{if } i \neq j. \end{cases}$$

The optimal policy will depends on δ, but not on x. Increases in the utility differential $u_G - u_B$ and reductions in learning costs both affect the optimal policy through increases in δ.

Substituting these payoffs in to the necessary and sufficient conditions (4) for some action $a_i \in B(P)$ yields

$$\frac{\delta \mu_i}{1 + \delta P(a_i)} + \sum_{a_j \in B(P)} \frac{\mu_j}{(1 + \delta P(a_j))} + \sum_{a_k \in A \setminus B(P)} \mu_k = 1$$

Since the latter two terms on the left-hand side are the same for all chosen acts, it follows that the optimal policy equalizes the first term:

$$\frac{\mu_i}{1 + \delta P(a_i)}$$

across $a_i \in B(P)$. This equality has two implications. First, $\mu_i > \mu_j$ implies $P(a_i)/P(a_j) > \mu_i/\mu_j$ so that actions that are more likely to be optimal are more than proportionately more likely to be taken. Second, if the first K actions are taken then the first order condition for the Kth action and the equality of the $\frac{\mu_i}{1+\delta P(a_i)}$ together imply

$$(\delta + K)\frac{\mu_K}{1+\delta P(a_K)} = \sum_{k=1}^{K} \mu_k$$

which requires μ_K to be greater than $\frac{1}{K+\delta}\sum_{k=1}^{K}\mu_k$. Suppose, for example, $\delta = 1$, then $\mu_2/(\mu_1+\mu_2)$ must be greater than $1/3$ if the first two actions are to be considered, $\mu_5/(\sum_1^5 \mu_a)$ must be greater than $1/6$ if the first five actions are to be considered, and so on.

Theorem 1 provides a complete characterization of the optimal strategy for this problem according to the Shannon Model.

Theorem 1 *If $\mu_M > \frac{1}{M+\delta}$ define $K = M$. If $\mu_M < \frac{1}{M+\delta}$, then define $K < M$ as the unique integer such that,*

$$\mu_K > \frac{\sum_{k=1}^{K}\mu(\omega_k)}{K+\delta} \geq \mu_{K+1}. \tag{8}$$

Then the optimal attention strategy involves,

$$P(a_i) = \frac{\mu(\omega_i)(K+\delta) - \sum_{k=1}^{K}\mu(\omega_k)}{\delta\sum_{k=1}^{K}\mu(\omega_k)} > 0 \tag{9}$$

for $i \leq K$, with $P(a_i) = 0$ for $i > K$. Furthermore the posteriors associated with all chosen options $a_i \leq K$ take the same form,

$$\gamma^i(\omega_j) = \begin{cases} \frac{(1+\delta)\sum_{k=1}^{K}\mu(\omega_k)}{K+\delta} \text{ for } i = j; \\ \frac{\sum_{k=1}^{K}\mu(\omega_k)}{K+\delta} \text{ for } i \neq j \text{ and } j \leq K \\ \mu(\omega_j) \text{ for } j > K. \end{cases}$$

This solution has two striking features. The first is that many alternatives are never chosen - in particular those for which the prior probability that the good is of high quality is low. Moreover, nothing is learned about these alternatives: the posterior probability that these goods are of high quality is the same as the prior probability, regardless of which good is actually chosen. This highlights the connection between the Shannon model and the theory of consideration sets. However, the second striking feature is that, unlike in standard consideration set models, 'mistakes' are still made in choice between considered alternatives: the probability that the high quality good is chosen is below one, even if this good is in the consideration set.

A numerical example highlights these features.

3.1.2 Example 1

Suppose that the $u_G = 1$ and $u_B = 0$, there are 10 possible alternatives, and prior beliefs are distributed exponentially according to $\mu(\omega_k) = \alpha\beta^{k-1}$, for $1 \leq k \leq 10$, with $\alpha = \frac{1-\beta}{1-\beta^{10}}$ to ensure that the prior probabilities sum to 1. In this case the number K of chosen options satisfies,

$$(K+\delta)\,(1-\beta)\,\beta^{K-1} > \left(1-\beta^{K}\right) \geq (K+\delta)\,(1-\beta)\,\beta^{K}.$$

The precise nature of the solution will depend on the cost of information λ and the parameter of the exponential distribution β. Figure 1 illustrates the optimal policy for $\beta = 0.8$ and four different levels of λ. 'Prior' refers to the prior probability that each of the 10 alternatives is the good option. 'Prob chosen' is the (unconditional) probability that each alternative is chosen, while 'Posterior' is the posterior probability of each alternative being of high quality conditional on it being chosen.

Figure 1: Optimal Behavior in Example 1

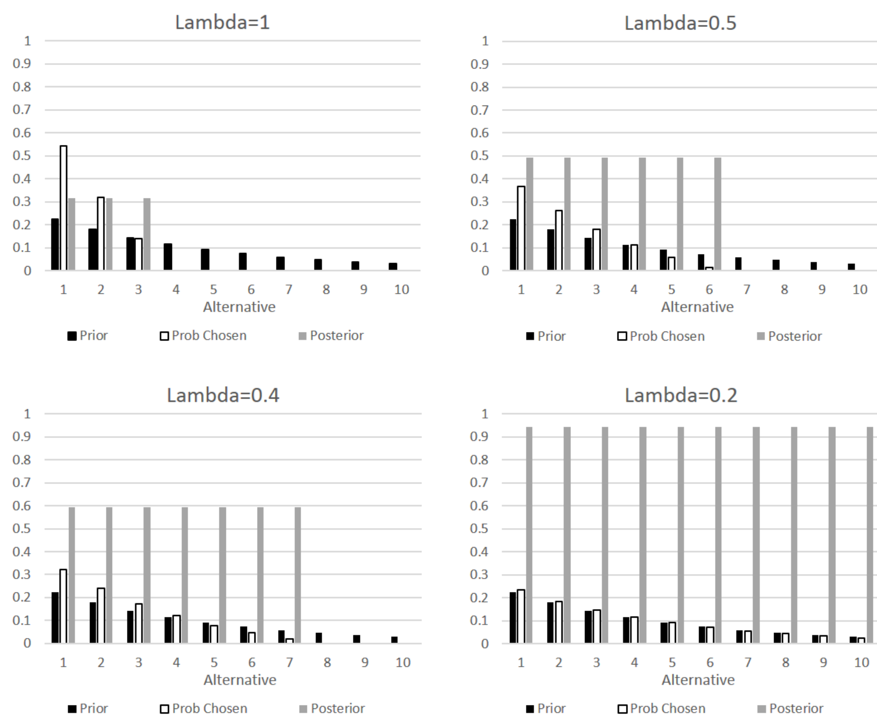

Figure 1 highlights the key features of the optimal policy according to the Shannon model in this setting. Looking first at the top left hand panel ($\lambda = 1$) we see that there is a distinct 'consideration set' of three items which are chosen with positive probability. None of the other alternatives are ever chosen. However, even within the consideration set, the subject makes 'mistakes' - for each of the chosen alternatives, the probability of it in fact being high quality is about 31%. Strikingly, this figure is exactly the same for all chosen alternatives: despite the fact they had different prior probabilities of being high quality, the decision maker learns exactly enough to make them all identical ex post, conditional on being chosen. We return to this observation in section 3.1.4 below. Finally, as λ decreases, the size of the consideration set increases and the probability of a mistaken choice falls, as can be seen from panels 2-4.

3.1.3 Comparison with Necessary Conditions

This application highlights the advantage of the necessary and sufficient conditions described in Proposition 1 over the necessary conditions introduced in MM. The key observation is that the MM conditions do not make any reference to the unchosen actions, and do not therefore provide a corresponding check on whether or not higher expected utility would be available were the set of chosen actions were to be changed. The following, which is an immediate corollary of Theorem 1, identifies all sets $C \subset A$ such that there exists a solution to the MM necessary conditions that specifies strictly positive probabilities for all options in C.

Corollary 1 *Consider any non-empty set of options C, and re-index options so that the first $|C|$ indices specify the probabilities of all options in C, retaining the diminishing order $\mu(\omega_i) \geq \mu(\omega_j)$ all $1 \leq i \leq C-1$. Then there exists a solution to the MM necessary conditions with all probabilities $P(a) > 0$ for all $a \in C$ if and only if,*

$$\frac{\mu(\omega_{|C|})}{\sum_{k=1}^{C} \mu(\omega_k)} > \frac{1}{|C| + \delta}.$$

Note that singleton sets always satisfy this condition. Note also that all subsets of the true optimal set as defined by Theorem 1 also satisfy these conditions. How many other sets satisfy them depends on model parameters. As δ increases, so ever more sets satisfy the conditions. In the limit, for δ so high that the true optimum is to pick all options with strictly positive probability, all subsets of available options satisfy the conditions of the corollary.

3.1.4 Relation to the Drift Diffusion Model

A further striking feature of the solution to the above problem is that posterior beliefs associated with chosen actions do not depend on prior beliefs, but only on relative payoffs. This is apparent from the Invariant Likelihood Ratio property introduced in Proposition 2.

This means that the probability of each alternative being of high quality *conditional on being chosen* is identical and so independent of prior beliefs.

This highlights a similarity between the solution of the Shannon model and the Drift Diffusion Model (DDM) that has been heavily studied in psychology (see for example Ratcliff [1978], Busemeyer and Townsend [1993]). In this model, a DM faced with a choice between two alternatives is assumed to sequentially receive information about the value of each option. The resulting beliefs follow a diffusion process, the drift of which is determined by the true quality of each. It is further assumed that information accumulation continues until the probability that one alternative is of high quality surpasses some boundary, at which point that alternative is chosen. In the basic version of the DDM, these boundaries are time invariant. Recent theoretical work (Fudenberg *et al.* [2015]) shows that such time invariant boundaries are optimal in some circumstances.

Our result highlights a link between the Shannon model and the DDM. In both cases, the posterior probability that an option is of high quality conditional choice is always the same, regardless of which alternative is chosen. In the DDM this stems directly from the time invariant boundaries that trigger choice, while in the case of the Shannon model it derives from the ILR property of proposition 2.

This DDM has been applied to situations very similar to the one described in section 3.1: A subject (typically a monkey) is presented with evidence indicative of which of several options will yield a known prize (typically juice). The evidence is costly to process, so that decisions take time, and mistakes are made. One key subject of study is the pattern of the resulting mistakes and how they relate to the quality of the evidence, the number of available options, prior beliefs, the nature of the prize, and the time taken to decide. This experimental design has become the archetype for exploration of the perceptual process and its neurological correlates in decision tasks. The two option version of this task is particularly well-studied, and the DDM has gained centrality due to a certain level of behavioral and even neurophysiological support. The model also has proponents within economics, (for example Krajbich and Rangel [2011]), who are developing multi-option generalizations under exploration.

3.2 Consumer Problem 2: Independent Valuations

We now consider the case in which the consumer is faced with the choice between a number of different alternatives, the values of which are uncertain, but independently distributed.

For example, a decision maker may be choosing which of several different cars to buy. Each car has a distribution of possible utilities it can deliver, depending on its price, fuel efficiency, reliability and so on. We make the assumption that the utility associated with one car is independent of that of any other. The consumer must decide which cars to consider, and what to learn about each considered car prior to purchase. The assumption of independence means that learning about the quality of one car does not imply anything about any other car.

The consumer is again faced with the choice of M possible actions $A = \{a_1, .., a_M\}$. Let $X \subset \mathbb{R}$ be the (finite) set of possible utility levels for all acts. We define the state space as $\Omega = X^M$. A typical state is therefore a vector of realized utilities for each possible action:

$$\omega = \begin{pmatrix} \omega_1 \\ \omega_2 \\ \vdots \\ \omega_M \end{pmatrix}$$

where $\omega_i \in X$ for all $a_i \in A$. The utility of a state/action pair is then given by

$$u(a_i, \omega) = \omega_i$$

The assumption of independence implies that there exists probability distributions $\mu_1, .., \mu_M \in \Delta(X)$ such that, for every $\omega \in \Omega$:

$$\mu(\omega) = \Pi_{i=1}^{M} \mu_i(\omega_i)$$

We call a decision problem with this set up an *independent consumption problem.*

The optimal approach to information acquisition in the independent consumption problem once again includes a cut-off strategy determining a consideration set of alternatives about which the consumer will learn and from which they will make their eventual choice. However in this case, the cut-off is in terms of the expectation of the normalized utilities $z(a, \omega) \equiv \exp(u(a, \omega)/\lambda)$ evaluated at prior beliefs.

Theorem 2 *Any optimal policy for an independent consumption problem will have a cut-off $c \in \mathbb{R}$ such that, for any $a_i \in A$, $P(a_i) > 0$ if*

$$\sum_{\omega \in \Omega} z(a_i, \omega)\mu(\omega) = \sum_{\omega \in \Omega} \exp(\omega_i/\lambda)\mu_i(\omega_i) > c$$

and $P(a_i) = 0$ *otherwise*

Like the 'find the best alternative' case, we can think of the consumer as ranking alternatives and including the best alternatives in the consideration set. Here the ranking depends on the expectation of the transformed net utilities $z(a_i, \omega)$ at the prior beliefs. This transformation reflects the importance of information acquisition. Consider two actions with the same ex ante expected payoffs, a safe action that pays its expected value in every state and a risky one whose payoff varies across states. In this case, the second risky action will be more valuable, since the decision maker can tailor their information strategy in such a way that they take this action in high valuation states and avoid this action in low valuation states. This explains the convex transformation of the payoffs: variance is valuable in a learning environment.

This also explains the role of the information cost. As λ rises, information becomes more costly the ability to tailor choice to the state disappears. As λ approaches infinity, $Ez(a_i, \omega)/Ez(a_j, \omega)$ approaches $Eu(a_i, \omega)/Eu(a_j, \omega)$ and choice is based on ex ante expected payoffs. As λ approaches zero, information becomes free and the best action is chosen in each state. An action remains unchosen only if its maximal payoff is less than the minimum payoff to some other action. In this case $Ez(a_i, \omega)/Ez(a_j, \omega)$ approaches infinity for all j if and only if a_i is chosen.

In the 'independent' valuation case, the ordering of actions is far from obvious without applying the necessarily and sufficient condition. As the next example illustrates this means that the nature of the consideration set can change in surprising, and non-monotonic ways with the cost of attention.

3.2.1 Example 2

Consider an independent consumption problem in which there are three possible utility levels,with $X = \{0, 5.5, 10\}$. There are six available actions. The first (a_1) has a value of 5.5 for sure and so is the 'safe' option. The other five $(a_2, ..., a_6)$ have an ex-ante 50% chance of having value 10 and a 50% having value zero, and so can be seen as 'risky' options.

In order to characterize behavior in this setting, the following lemma is of use

Lemma 1 *Let* (μ, A) *be an independent consumption problem and* $\{a_1, ...a_N\} = B \subset A$ *be a set of ex ante identical actions (i.e.* $\mu_i(x) = \mu_j(x) = \mu_B(x)$ *for all* $x \in X$ *and* $i, j \leq N$*).*

Then a strategy that picks each $a_i \in B$ with the unconditional probability $\frac{1}{N}$ and assigns unconditional probabilities according to equation 2 is optimal if, for each $a_j \notin B$

$$\sum_{x \in X} \exp(x/\lambda)\mu_j(x) \leq \frac{1}{n}\left[\sum_{\bar{x} \in X^N} \frac{\Pi_{n=1}^{N}\mu_B(\bar{x}_n)}{\sum_{n=1}^{N} \exp\left(\bar{x}_n/\lambda\right)}\right]^{-1}$$

The Shannon model implies that there are 3 possible consideration sets for this problem: only a_1 (safe only), only $a_2 - a_5$ (risky only), or all options (safe and risky options). Each of these is the optimal consideration set at different levels of attention costs. Perhaps surprisingly, the relationship between attention costs and the associated consideration set is non-monotonic. For example, for $\lambda = 30$ only the safe option is chosen with positive probability, for $\lambda = 20$ both the safe and risky options will be used, for $\lambda = 2$ only the risky options will be used, while for $\lambda = 1$ again all options are used. Figure 2 shows the unconditional probability of the sure thing and each of the risky alternatives being chosen at each value of λ.

Figure 2: Unconditional Choice Probabilities in Example 2.[9]

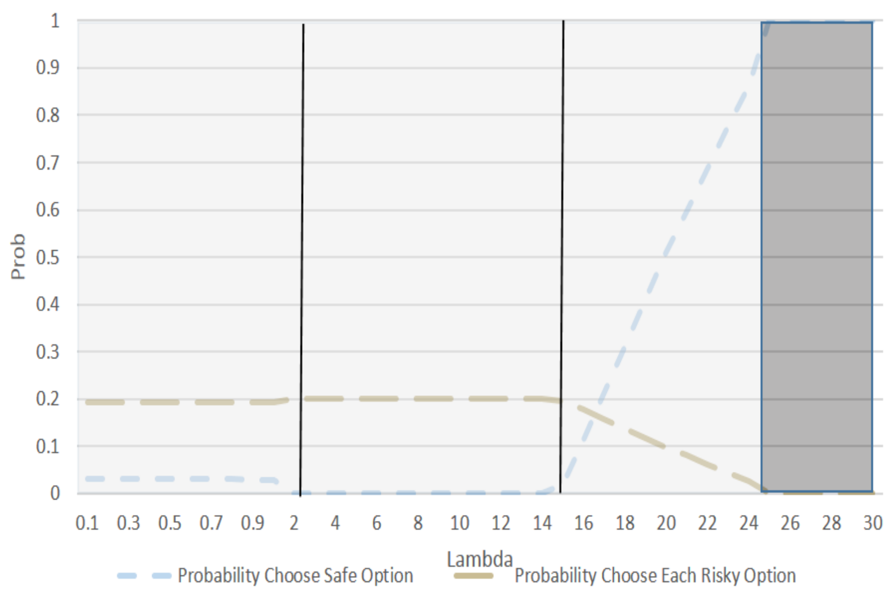

3

Figure 2 clearly shows 4 different regions for the parameter λ, each of which is related to a different consideration set. For very low values of λ, when information is very cheap, all alternatives are used. As λ increases, the probability that the safe option is chosen drops to zero. For still higher values of λ, the sure thing is once again used, along with the risky options. For the highest values of λ (when information is very expensive) only the sure thing

[9]Note that the scale uses increments of 0.1 for $\lambda < 1$, and 1 for $\lambda > 1$

is used.

How is this change occurring, given the condition of Theorem 2? It turns out that increasing the value of λ has two distinct effects on behavior. First, it can change the ranking of the the risky and safe options in terms of their normalized utilities given prior beliefs. At low levels of λ, the risky option dominates the sure thing (the light grey region in Figure 2). However, as attention costs rise, eventually the normalized utility of the sure thing moves above that of the risky option (the dark grey region). Thus, for low attention costs, if only one option is to be used it must be the risky option, while for high costs, only the sure thing can be used on its own.

A second effect of rising attention costs is to change whether all alternatives are above the threshold, or just some subset of them: At very low cost levels, all options are above the threshold. As costs increase, the sure thing drops below the threshold and only the risky options are used. Further cost rises lead to the sure thing moving back above the cost threshold.

Some intuition for this effect can be gained from Figure 3. This shows the probability that the sure thing is chosen conditional on all the risky options being bad, and the probability that all the risky options are bad conditional on the sure thing being chosen. It illustrates that the sure thing plays very different roles in the consideration set at low and high information cost levels. At low cost levels, when the consumer is very well informed, the sure thing is only chosen when it is known with high probability that all the risky options are of low quality: in other words it is only chosen when it is actually the best option. As information costs rise, at some point it becomes too costly for the consumer to identify such states of the world, so the sure thing is no longer used. When the sure thing again enters the consideration set at higher cost levels, it is used in a very uninformed manner. The probability that all the risky options are bad if the sure thing is chosen is only about 4%, or slightly above the prior belief that this is the case. In this part of the attention cost region, the sure thing is used by the consumer as a way of mixing in an uninformed choice with their informed choice in order to lower costs. As λ rises, use of this unformed option rises until eventually use of the risky options ceases.

Figure 3: Conditional Probabilities in Example 2.[10]

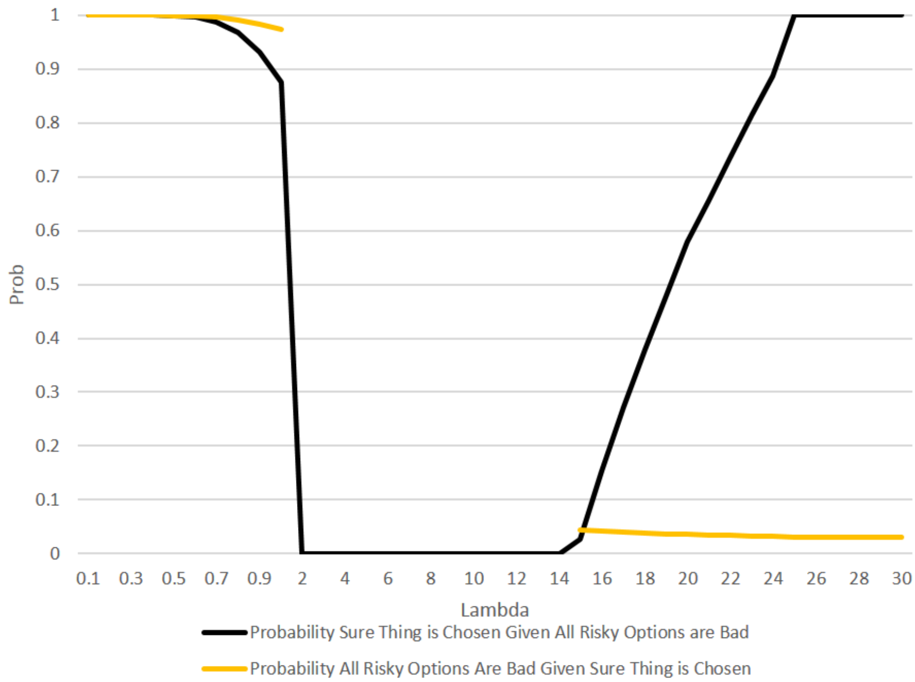

3.2.2 Information Cost and Risk Aversion

Our conditions can also be used to explore the relationship between risk aversion and information costs in consideration set formation. Intuitively, such a relationship exists because a more risk-averse individual requires a higher degree of certainty before they are prepared to choose a risky option, and that higher degree of certainty requires more information. Thus, an investor may be prepared to invest in risky stocks if the have low risk aversion, or low information costs, but not if they are risk averse and find information costly.

In order to explore this trade off, we can modify the set up of Example 2, so that the payoffs of each alternative are denominated in monetary units. The risky options pay off \$10 and \$0 if they are good or bad, while the sure thing pays of \$5 for sure. Monetary payoffs are converted to utility using the function

$$u(x) = \frac{x^{1-\rho}}{1-\rho}$$

[10]Note that the scale uses increments of 0.1 for $\lambda < 1$, and 1 for $\lambda > 1$

The parameter ρ determines the degree of risk aversion, with $\rho = 0$ equivalent to risk neutrality.

We can now use Lemma 1 to map out the optimal consideration sets as a function of λ and ρ. The results are shown in Figure 4, which demonstrates a surprisingly complex relationship between the two parameters. Broadly speaking, the intuition described above holds: use of the risky options increases both with lower information costs and lower risk aversion. At very low information costs, all options appear in the consideration set. However, at higher information costs, there are still values for ρ for which all options are chosen with positive probability. This occurs at intermediate levels of risk aversion, between the 'risky asset only' and 'safe asset only' areas of the parameter space. This region corresponds to the 'uninformed' use of the safe option described above.

Figure 4: Consideration Sets as a Function of ρ and λ

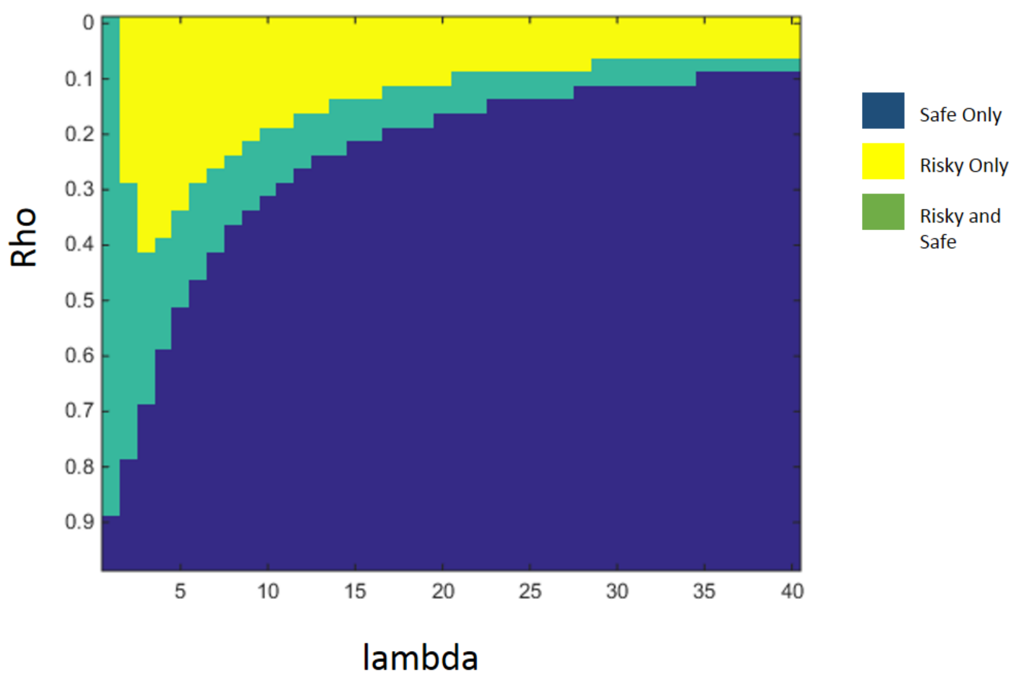

3.3 Consumer Problem 3: Correlated Valuation

The general case in which the value of the different alternatives may be correlated is more complex. We begin with a geometric interpretation of the model which borrows from Caplin *et al.* [2016]. We then present a simple 'market entry' test based on Proposition 1. We conclude this section with an example.

3.3.1 A Geometric Interpretation

Bayes' rule implies a tight relationship between state dependent stochastic choice and posterior beliefs, $\gamma^a(\omega) = P(a|\omega)\mu(\omega)/P(a)$. We can use this relationship to rewrite problem (1) replacing $P(a|\omega)$ with $\gamma^a(\omega)$ and $P(a)$. The resulting maximization problem is of the form:

$$\begin{aligned} V(\mu, A) \quad = \quad & \max_{\{\{P(a)\}_{a\in A}, \{\gamma^a\}_{a\in B(P)}\}} \sum_{a\in B(P)} P(a) \sum_{\omega\in\Omega} \gamma^a(\omega) u(a,\omega) \\ & -\lambda \left[\sum_{a\in A} P(a) \sum_{\omega\in\Omega} \gamma^a(\omega) \ln \gamma^a(\omega) - \sum_{\omega\in\Omega} \mu(\omega) \ln \mu(\omega) \right] \end{aligned}$$

Let $N(\gamma^a) = \sum_{\omega\in\Omega} \gamma^a(\omega) u(a,\omega) - \lambda\gamma^a(\omega) \ln \gamma^a(\omega)$ denote the net utility of choosing action a and the resulting posterior γ^a, then

$$V(\mu, A) = \max_{\{\{P(a)\}_{a\in A}, \{\gamma^a\}_{a\in B(P)}\}} \sum_{a\in B(P)} P(a) N(\gamma^a) + \lambda \sum_{\omega\in\Omega} \mu(\omega) \ln \mu(\omega)$$

where the latter term is independent of optimization. The optimal set of posteriors is the one that maximizes the expected value of these net utilities.

Figure 5

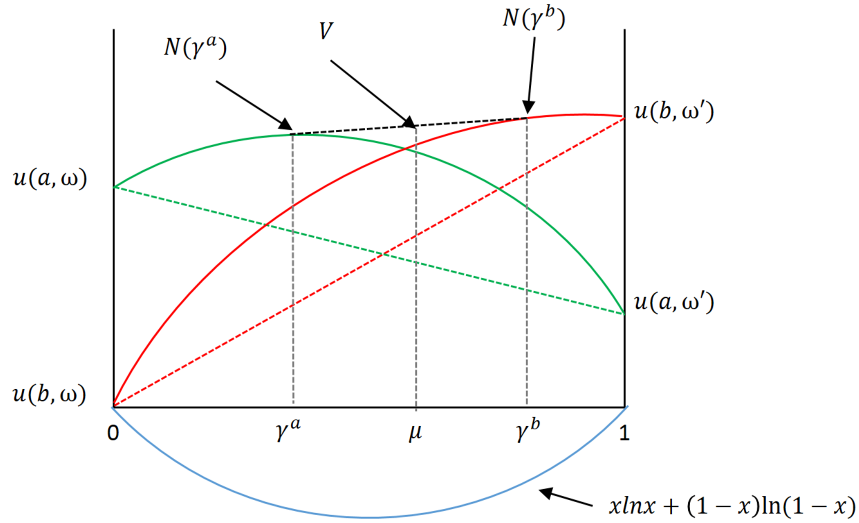

Figure 5 illustrates how this works in the simple case with two states $\{\omega, \omega'\}$ and two actions $A = \{a, b\}$. The probability of state ω' is represented on the horizontal axis. The cost associated with any posterior, $\lambda \sum_\omega \gamma(\omega) \ln \gamma(\omega)$, is given by the convex line below the axis. The expected payoff from action a as a function of the resulting posterior γ^a, $\sum_\omega \gamma^a(\omega) u(a, \omega)$, is given by the downward sloping dotted green line, while the upward sloping dotted red line illustrates the expected payoff to action b as a function of γ^b. The figure illustrates the case in which a pays more in state ω and b pays more in state ω'. The net utilities associated with each action are obtained by subtracting the information cost from the expected payoffs; these are represented by the solid green and red lines.

Given the prior μ, the value of a strategy that assigns probabilities to two posteriors γ^a and γ^b can be determined as follows. The net utilities associated with γ^a and γ^b are $N(\gamma^a)$ and $N(\gamma^b)$ respectively. Rational expectations implies that $P(a)\gamma^a + (1 - P(a))\gamma^b = \mu$. Hence the value of the strategy $P(a)N(\gamma^a) + (1 - P(a))N(\gamma^b)$ is a weighted average of these net utilities. In fact, it is the height above μ of the cord connecting $N(\gamma^a)$ and $N(\gamma^b)$. The optimal strategy can be found by identifying the posteriors that support the highest possible chord as it passes over the prior. The posteriors in the figure have this property, and so would form part of an optimal strategy for this decision problem. Caplin *et al.* [2016] show that this insight generalizes to problems with many actions and many states: one graphs the net utilities and finds the point on the convex hull directly above the prior; the optimal posteriors are the points of tangency of the supporting hyperplane at this point and the net utility functions.

Our interest in this paper is the implications of this characterization for the optimal consideration set. In order for an action to be considered, its net utility must touch the supporting hyperplane. Except in cases of indifference, this means that the net utility function associated with this action would pierce the hyperplane associated with a problem that did not include this act. This is clearly more likely if the net utility associated with this action is higher (i.e. the payoffs are higher) or the plane is lower (i.e. the payoffs to the other actions are lower). But it also matters in which states the action pays off relative to the slope of the hyperplane: if the action pays off in states in which the hyperplane is low, then it will more likely prove valuable.

To see this, Figure 6 illustrates the net utility function for a problem with one act. Since there is only one action the convex hull of net utility is net utility itself, and the point of tangency of the supporting plane lies directly above the prior so that the prior is the optimal posterior. This makes sense since there is no incentive to gather information. The black line illustrates the supporting plane. A second action will prove valuable if its net utility pierces this plane, so that the supporting plane to the new problem at μ is higher. When is this the case? As stated above, clearly an action that pays off more will have a greater chance. But the new action need not pay off more in all states. actions tend to be more valuable if they pay off in states in which the hyperplane is low. There are two reasons that this may be the case. First, the action in the figure pays off more in state ω, which implies that net utility tends to slope downward and so the plane tends to be lower in state ω'. Actions that pay off in state ω' are therefore more likely to be valuable. This is a hedging motive. Second, actions that pay off in more likely states are more valuable. In the figure, μ places more weight on ω'. This shifts the supporting plane to a point with greater downward slope so that the plane is lower in ω'. Again actions that pay off in state ω' are more likely to be valuable.

Figure 6

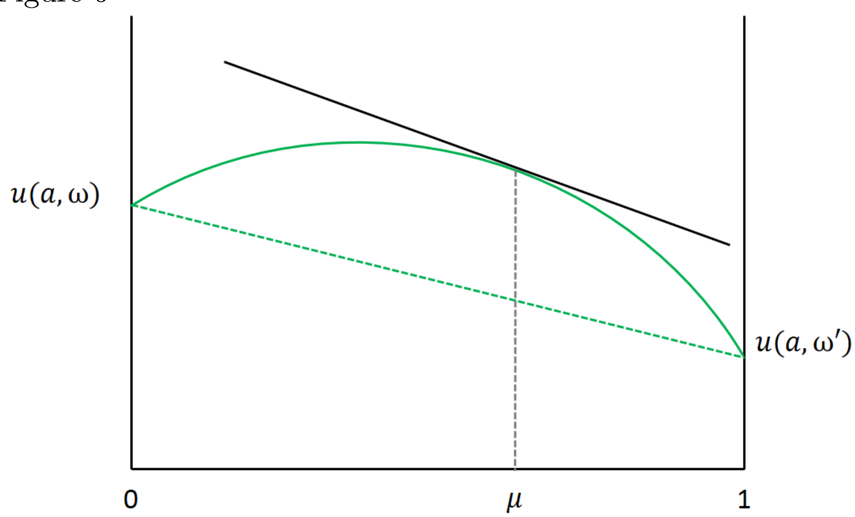

3.3.2 A Simple 'Market Entry' Test

Calculating whether a net utility function lies above or below a supporting hyperplane can be quite involved. Fortunately, the necessary and sufficient conditions provide a simple test for whether an action should be added to a consideration set. All unchosen actions $a \in A \backslash B(P)$ must satisfy the inequality

$$\sum_{\omega \in \Omega} \frac{z(a,\omega)\mu(\omega)}{\sum_{b \in A} P(b) z(b,\omega)} < 1$$

We can therefore solve the model without action a and check to see whether this inequality is satisfied for a. This approach might prove particularly useful in models of market entry in which the $P(b)$ represent the equilibrium pre-entry.

What types of goods pass the entry test?

$$\sum_{\omega\in\Omega}\frac{z(a,\omega)\mu(\omega)}{\sum_{b\in A}P(b)z(b,\omega)} = Ez(a,\omega) + E\frac{1}{\sum_{b\in A}P(b)z(b,\omega)} + \text{cov}\left(z(a,\omega),\frac{1}{\sum_{b\in A}P(b)z(b,\omega)}\right)$$

The first term is familiar from the case of uncorrelated acts: a higher expected $z(a,\omega)$ makes an action more desirable. The second term relates to the unconditional expected value of already chosen acts. This term is the same for all actions and therefore does not itself distinguish between chosen and unchosen acts. The third term is new and represents the hedging motive discussed above. A high covariance term means the action tends to pay off more in states in which other actions pay off less.

3.3.3 Example 3

Consider a choice set consisting of three alternatives, a, b and c. The value of each of these alternatives is determined by an underlying state drawn from $\Omega = \{\omega_1, \omega_2\}$, each of which is equally likely. The payoff of each action in each state (in utility terms) is described in Table 1.

Table 1

Alternative	ω_1	ω_2	$E(z(x))$
a	5	5	1.65
b	6	0	1.41
c	0	15	2.74

Using Proposition 1, it is easy to show that, for $\lambda = 10$, only alteratives b and c will be in the consideration set - i.e. will be chosen with positive probability. This is despite the fact that a has a higher expected normalized utility than b at prior beliefs, as can be seen from the last column of Table 1. The reason for this is that option b provides a better hedge than a for option c. The presence of option c induces the consumer to find out with high precision whether the state of the world is ω_1 or ω_2. Having done so, they sometimes learn that ω_1 is very likely to be the true state, in which case they prefer b to a.[11]

Example 3 illustrates that the cutoff strategy breaks down because the optimal strategy now potentially depends on all the available actions. Even risk neutral consumers may utilize the 'hedging' value of a given act, if it is of high quality in states where others are low quality. Such actions increase the value to learning, because it means appropriate action can be taken

[11] A similar example appears in Matejka and McKay [2015].

regardless of what is learned.

In the decision making environment of Example 3, a new alternative will not enter the consideration set unless it satisfies

$$\begin{aligned} 1 &\leq \frac{1}{2}\frac{z(d,\omega_1)}{P(b)z(b,\omega_1)+P(c)z(c,\omega_1)}+\frac{1}{2}\frac{z(d,\omega_2)}{P(b)z(b,\omega_2)+P(c)z(c,\omega_2)} \\ &= \frac{1}{2}\left[\frac{z(d,\omega_1)}{1.41}+\frac{z(d,\omega_2)}{2.74}\right] \end{aligned}$$

This condition can then be used to find the 'minimum cost' way of ensuring that a product will be enter into the consideration set. In other words, the assignment of $u(d,\omega_1)$, $u(d,\omega_2) \geq 0$ which guarantees that d be in the consideration set, while minimizing the expected utility of d at prior beliefs. In the above example, it is clear that the solution to this problem is to set $u(d,\omega_2) = 0$ and set $u(d,\omega_1)$ in order to make $\frac{1}{2}\frac{z(d,\omega_1)}{1.41} = 1$. This allocation puts maximal utility on the state which has the lowest expected value of normalized utility given current choice patterns.

4 Literature Review

Our paper provides new techniques for solving models or rational inattention which have been popular in economics since their introduction by Sims [2003].[12] Specifically, we augment the results of Matejka and McKay [2015] to provide conditions which are both necessary and sufficient for optimality.

Within economics, there are several recent papers that have tackled the concept of consideration sets from a theoretical perspective. Masatlioglu *et al.* [2012] (henceforth MNO) take a 'revealed preference' approach, using the identifying assumption that if an alternative x is not in the consideration set for some choice set S, removing x will not change the consideration set. They use this condition to provide necessary and sufficient conditions for a data set to be consistent with choice from consideration sets.[13] Unlike our approach, consideration sets in MNO do not come about as the result of optimizing behavior, although our model does satisfy their identifying restriction. This means on the one hand that their model is potentially more flexible, while on the other it provides fewer comparative static

[12] See for example the application of the model to investment decisions (e.g van Nieuwerburgh and Veldkamp [2008]), global games (Yang [2015]), and pricing decisions (Mackowiak and Wiederholt [2009], Matějka [2015], Martin [2013]).

[13] See Lleras *et al.* [2010] and Dean *et al.* [2015] for similar approaches.

predictions. MNO also assume the absence of mistakes within the consideration set. A combination of deterministic consideration sets and preference maximization mean that MNO's model predict that choice will be deterministic. More recently, Demuynck and Seel [2014] have also taken a revealed preference approach to consideration set formation, in which they assume that some commodities in a bundle may not be observed. Again, the resulting model does not have an optimizing component, and choice is deterministic.

Another recent approach is that of Manzini and Mariotti [2014] (henceforth MM), who assume that consideration sets are formed stochastically, with any given alternative having a fixed probability of being considered. Again, consideration is not the result of optimization, and choice within the consideration set is always optimal, but the random nature of consideration leads to random choice. Unlike our model, all alternatives are always chosen, meaning that, by our definition, all alternatives are in the consideration set. On a technical level, our model would violate the I-Asymmetry and I-independence axioms of MM.

A further model of consideration set formation is the search and satisficing approach, originally suggested by Simon [1955]. In such a model, alternatives are considered one by one until one is found which is 'good enough'. As discussed in Caplin *et al.* [2011], satisficing can be seen as the result of an optimizing procedure in the face of information costs. While this paper considers only the case in which alternatives are ex ante identical, Gabaix *et al.* [2006] discuss the extension to a situation in which the decision maker may have different priors about the quality of different alternatives. However, the same paper shows that extending this model to the case in which search only reveals partial information about the quality of the alternative is generally intractable.

Their is a much longer history of research into consideration sets in marketing. Classic examples include Roberts and Lattin [1991] and Hauser and Wernerfelt [1990]. Typically, these papers run into the same tractability problems discussed in Gabaix *et al.* [2006]. These are solved by making strong distributional assumptions on the nature of information, and using this to derive various moments of the choice distribution. As far as we are aware, none of these papers have considered the approach of using rational inattention to model consideration set formation.

5 Conclusion

This paper makes two distinct contributions. First, we introduce new necessary and sufficient conditions for the solution of the Shannon model. These conditions are an important advance for anyone interested in applying the model, as they allow for the identification of the set of actions that will be chosen with positive probability at the optimum. Without identifying this set, the necessary conditions of MM are not useful for characterizing behavior. We hope that these conditions will aid the adoption of the Shannon model as a benchmark way of incorporating the insights of rational inattention into economic analysis.

Second, we show how to use the Shannon model as a model of optimal consideration set formation. This approach has a number of advantages over most existing approaches - allowing us to identify the optimal consideration set, and the optimal pattern of mistakes within the consideration set in a tractable way. We hope that this application will provide new tools to researchers interested in marketing and demand system estimation, where the identification of the set of considered options is of significant importance.

References

Marina Agranov and Pietro Ortoleva. Stochastic choice and preferences for randomization. *Available at SSRN 2644784*, 2015.

Jose Apesteguia and Miguel A Ballester. Single-crossing random utility models. 2016.

Jerome R Busemeyer and James T Townsend. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, 100(3):432, 1993.

Andrew Caplin and Mark Dean. Behavioral implications of rational inattention with shannon entropy. NBER Working Papers 19318, National Bureau of Economic Research, Inc, August 2013.

Andrew Caplin, Mark Dean, and Daniel Martin. Search and satisficing. *American Economic Review*, 101(7):2899–2922, December 2011.

Andrew Caplin, Mark Dean, and John Leahy. Rationally Inattentive Behavior. 2016. mimeo.

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

Andrew Ching, Tülin Erdem, and Michael Keane. The price consideration model of brand choice. *Journal of Applied Econometrics*, 24(3):393–420, 2009.

T. Cover and J. Thomas. *Elements of Information Theory 2nd Edition*. John Wiley and Sons, Inc., New York, 2006.

Geoffroy De Clippel, Kfir Eliaz, and Kareen Rozen. Competing for consumer inattention. *Journal of Political Economy*, 122(6):1203–1234, 2014.

M. Dean, O. Kibris, and Y Masatlioglou. Limited Attention and Status Quo Bias. 2015. mimeo.

Thomas Demuynck and Christian Seel. Revealed preference with limited consideration. *Available at SSRN 2509353*, 2014.

Kfir Eliaz and Ran Spiegler. Consideration sets and competitive marketing. *The Review of Economic Studies*, 78(1):235–262, 2011.

Drew Fudenberg, Philipp Strack, and Tomasz Strzalecki. Stochastic choice and optimal sequential sampling. *Available at SSRN 2602927*, 2015.

Xavier Gabaix, David Laibson, Guillermo Moloche, and Stephen Weinberg. Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96(4):1043–1068, September 2006.

Sen Geng. Decision time, consideration time, and status quo bias. *Economic Inquiry*, 54(1):433–449, 2016.

John R. Hauser and Birger Wernerfelt. An evaluation cost model of consideration sets. *Journal of Consumer Research*, 16(4):pp. 393–408, 1990.

Wayne D. Hoyer. An examination of consumer decision making for a common repeat purchase product. *Journal of Consumer Research*, 11(3):pp. 822–829, 1984.

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

Nicola Lacetera, Devin G. Pope, and Justin R. Sydnor. Heuristic thinking and limited attention in the car market. *American Economic Review*, 102(5):2206–36, September 2012.

J. S. Lleras, Y Masatlioglou, D Nakajima, and E Ozbay. When More is Less: Limited Consideration. 2010. mimeo.

Bartosz Mackowiak and Mirko Wiederholt. Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803, June 2009.

Paola Manzini and Marco Mariotti. Stochastic choice and consideration sets. *Econometrica*, 82(3):1153–1176, 2014.

Paola Manzini and Marco Mariotti. Dual random utility maximisation. 2016.

Daniel Martin. Strategic pricing with rational inattention to quality. Mimeo, New York University, 2013.

Yusufcan Masatlioglu, Daisuke Nakajima, and Erkut Y. Ozbay. Revealed attention. *American Economic Review*, 102(5):2183–2205, August 2012.

Filip Matejka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.

Filip Matejka and Christopher A Sims. Discrete actions in information-constrained tracking problems. *CERGE-EI Working Paper Series*, (441), 2011.

Filip Matějka. Rationally inattentive seller: Sales and discrete pricing. *The Review of Economic Studies*, page rdv049, 2015.

Frederick Mosteller and Philip Nogee. An experimental measurement of utility. *The Journal of Political Economy*, pages 371–404, 1951.

Roger Ratcliff. A theory of memory retrieval. *Psychological review*, 85(2):59, 1978.

Elena Reutskaja, Rosemarie Nagel, Colin F Camerer, and Antonio Rangel. Search dynamics in consumer choice under time pressure: An eye-tracking study. *The American Economic Review*, 101(2):900–926, 2011.

John H. Roberts and James M. Lattin. Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 28(4):pp. 429–440, 1991.

Babur De Los Santos, Ali Hortacsu, and Matthijs R. Wildenbeest. Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, 102(6):2955–80, October 2012.

Herbert A Simon. A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 69(1):99–118, 1955.

Christopher A. Sims. Implications of Rational Inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.

Stijn van Nieuwerburgh and Laura Veldkamp. Information Immobility and the Home Bias Puzzle. *Journal of Finance (forthcoming)*, 2008.

Ming Yang. Coordination with flexible information acquisition. *Journal of Economic Theory*, 158:721–738, 2015.

6 Appendix

Proof of Proposition 1. Begin with the objective function from equation (1), and substitute in for the MM conditions for $P(a|\omega)$

$$\begin{aligned}
&\sum_{\omega\in\Omega}\sum_{a\in A}\mu(\omega)P(a|\omega)\left(u(a,\omega)-\lambda P(a|\omega)\right)+\sum_{a\in A}P(a)\ln P(a) \\
= &\sum_{\omega\in\Omega}\sum_{a\in A}\mu(\omega)P(a|\omega)\left(u(a,\omega)-\lambda\ln\left[\frac{P(a)z(a,\omega)}{\sum_{b\in A}P(b)z(b,\omega)}\right]\right)+\lambda\sum_{a\in A}P(a)\ln P(a)
\end{aligned}$$

we can rewrite the term in parentheses as

$$\begin{aligned}
&u(a,\omega)-\lambda\ln\left[\frac{P(a)z(a,\omega)}{\sum_{b\in A}P(b)z(b,\omega)}\right] \\
= \;&u(a,\omega)-\lambda\ln P(a)-\lambda\ln z(a,\omega)+\lambda\ln\sum_{b\in A}P(b)z(b,\omega) \\
= \;&-\lambda\left[\ln P(a)-\ln\sum_{b\in A}P(b)z(b,\omega)\right]
\end{aligned}$$

■

Substituting this back in to the objective function gives

$$-\lambda\sum_{\omega\in\Omega}\sum_{a\in A}\mu(\omega)P(a|\omega)\ln P(a)+\lambda\sum_{\omega\in\Omega}\sum_{a\in A}\mu(\omega)P(a|\omega)\ln\sum_{b\in A}P(b)z(b,\omega)+\lambda\sum_{a\in A}P(a)\ln P(a)$$

The first and last terms cancel out, and the logarithm in the middle term is not a function of a, leaving

$$\sum_{\omega\in\Omega}\lambda\ln\left(\sum_{a\in A}P(a)z(a,\omega)\right)\mu(\omega).$$

This new objective is concave in $P(a)$ and the constraints on $P(a)$ are linear. Hence the Kuhn-Tucker conditions are necessary and sufficient. The first order conditions for $P(a)$ is

$$\sum_{\omega\in\Omega}\lambda\frac{z(a,\omega)}{\sum_{b\in A}P(b)z(b,\omega)}\mu(\omega)+\varphi+\xi^a=0$$

The complementary slackness condition is $\xi^a P(a)=0$ and $\xi^a<0$.

If $P(a) > 0$ then

$$\sum_{\omega \in \Omega} \lambda \frac{z(a,\omega)}{\sum_{b \in A} P(b) z(b,\omega)} \mu(\omega) = -\varphi$$

Multiplying by $P(a)$ and summing over a

$$\lambda = -\varphi.$$

This holds as $P(a) > 0$ for some a. So if $P(a) > 0$

$$\sum_{\omega \in \Omega} \frac{z(a,\omega)}{\sum_{b \in A} P(b) z(b,\omega)} \mu(\omega) = 1$$

If $P(a) = 0$ then $\xi^a > 0$

$$\sum_{\omega \in \Omega} \frac{z(a,\omega)}{\sum_{b \in A} P(b) z(b,\omega)} \mu(\omega) = 1 - \frac{\xi^a}{\lambda} < 1$$

This completes the proof.

Proof of Theorem 1 . A necessary condition for chosen set $B \subset A$ to be optimal is that, for each $a \in B$,

$$\sum_{\omega \in \Omega} \frac{z(a,\omega)}{\sum_{b \in A} P(b) z(b,\omega)} \mu(\omega) = 1$$

Substituting in the specific payoffs, the key equation has simple form,

$$\begin{aligned} \sum_{\omega \in \Omega} \frac{z(a,\omega)}{\sum_{b \in A} P(b) z(b,\omega)} \mu(\omega) &= \frac{x(1+\delta)\mu_i}{x(1+\delta P(a_i))} + \sum_{a_j \in B \setminus a_i} \frac{x\mu_j}{x(1+\delta P(a_j))} + \sum_{a_k \in A \setminus B} \mu_k \\ &= \frac{\delta \mu_i}{1+\delta P(a_i)} + \sum_{a_j \in B} \frac{\mu_j}{(1+\delta P(a_j))} + \sum_{a_k \in A \setminus B} \mu_k = 1. \end{aligned}$$

In light of the fact that the sum of all priors is 1 we can subtract $\sum_{a_k \in A \setminus B} \mu_k$ from both sides and define the key ratio,

$$\rho_i = \frac{\mu_i}{1+\delta P(a_i)},$$

for all $a_i \in B$, we arrive at,

$$\delta \rho_i + \sum_{a_j \in B} \rho_j = \sum_{a_j \in B} \mu_j \Longrightarrow \rho_i = \frac{\sum_{a_j \in B} \mu_j - \sum_{a_j \in B} \rho_j}{\delta}.$$

Note that the RHS above is independent of a_i. Hence all ρ_j for $a_j \in B$ are identical, whereupon substitution yields,

$$\rho_i = \frac{\sum_{a_j \in B} \mu_j}{|B| + \delta}.$$

We can back out the implied probabilities for $a_i \in B$ as,

$$1 + \delta P(a_i) = \frac{\mu_i}{\rho i} = \frac{\mu_i \left(|B| + \delta\right)}{\sum_{a_k \in B} \mu_k};$$

so that,

$$P(a_i) = \frac{\mu_a \left(|B| + \delta\right)}{\delta \sum_{a_k \in B} \mu_k} - \frac{1}{\delta}.$$

We now apply analogous logic to unchosen options $a_i \in A \backslash B$ for which the corresponding necessary inequality is,

$$\sum_{\omega \in \Omega} \frac{z(a,\omega)}{\sum_{b \in A} P(b) z(b,\omega)} \mu(\omega) = \sum_{a_j \in B} \frac{\mu_j}{(1 + \delta P(a_j))} + \sum_{a_k \in A \backslash B} \mu_k + \delta \mu_i \leq 1,$$

Substituting the known value of ρ_j for $a_j \in B$ and subtracting $\sum_{a_j \in A \backslash B} \mu_j$ from both sides in light of the fact that the sum of all priors is 1 we arrive at,

$$\begin{aligned} \delta \mu_i &\leq \sum_{a_j \in B} \left[\mu_j - \frac{\mu_j}{(1 + \delta P(a_j))} \right] = \sum_{a_j \in B} \left[\mu_j - \rho_j \right] \\ &= \sum_{a_j \in B} \left[\mu_j - \frac{\sum_{a_j \in B} \mu_j}{|B| + \delta}. \right] = \frac{\delta \sum_{a_j \in B} \mu_j}{|B| + \delta}. \end{aligned}$$

We conclude that the necessary and sufficient conditions for optimality are satisfied by the specification of a set of chosen options $B \subset A$ if and only if the implied probabilities of chosen options $a \in B$ are all non-negative,

$$a_i \in B \Longrightarrow \mu_i > \frac{\sum_{k \in B} \mu_k}{(|B| + \delta)}.$$

while the corresponding inequality is satisfied by unchosen options ,

$$a_j \in A \backslash B \Longrightarrow \mu_j \leq \frac{\sum_{k \in B} \mu_k}{(|B| + \delta)}.$$

Note that if $\mu_M > \frac{1}{M+\delta}$, then setting $B = A$ verifies all conditions. Otherwise, if $\mu_M \leq \frac{1}{M+\delta}$, then we identify integer $K < M$ such that,

$$\mu_K > \frac{\sum_{k=1}^{K} \mu_k}{K+\delta} \geq \mu_{K+1}.$$

To see that there is such an integer, note first that $\mu_1 > \frac{\mu_1}{1+\delta}$, and that,

$$\mu_M \leq \frac{\sum_{k=1}^{K} \mu_k}{M+\delta} = \frac{1}{M+\delta}.$$

Hence one can count up from $k = 1$ to identify the smallest integer $K < M$ such that,

$$\mu_K > \frac{\sum_{k=1}^{K} \mu_k}{K+\delta} \geq \mu_{K+1}.$$

With such K identified, the fact that $\mu_K > \frac{\sum_{k=1}^{K} \mu_k}{K+\delta}$ verifies that the same is true for all $a_i \in A$ for which $\mu_i \geq \mu_K$, while the fact that $\mu_{K+1} \leq \frac{\sum_{k=1}^{K} \mu_k}{K+\delta}$ verifies that the same is true for all $a_j \in A$ for which $\mu_j \leq \mu_K$. Hence we have satisfied all necessary and sufficient conditions to identify action set $B = \{1, 2, .., K\}$ and corresponding probabilities $P(a_i) > 0$ on $a_i \in B$ to characterize the optimal attention strategy.

Note that for $a_i \in B$ the posteriors satisfy,

$$\gamma^i(\omega_j) = \frac{z(a_{i,}, \omega_j)\mu_j}{\sum_{k=1}^{K} P(a_k) z(a_k, \omega_j)}.$$

Hence, for $b = a$,

$$\gamma^i(\omega_i) = \frac{x(1+\delta)\mu_i}{\sum_{k=1}^{K} P(a_k) z(a_k, \omega_j)} = \frac{x(1+\delta)\mu_i}{x(1+\delta P(a_i))} = (1+\delta)\rho_i = \frac{(1+\delta)\sum_k \mu_k}{K+\delta}.$$

Moreover for $b \neq a$ and $b \leq K$,

$$\gamma^i(\omega_j) = \frac{x\mu_j}{\sum_{k=1}^{K} P(a_k) z(a_k, \omega_j)} = \frac{x\mu_j}{x(1+\delta P(a_i))} = \rho_i = \frac{\sum_k \mu_k}{K+\delta}.$$

Finally, for $b > K$,

$$\gamma^i(\omega_j) = \frac{x\mu j}{\sum_{k=1}^{K} P(a_k) z(a_k, \omega_j)} = \frac{x\mu_i}{x} = \mu_j.$$

This completes the proof. ■

Proof of Theorem 2 The proof is by contradiction. the expected utility of an action a_i at prior beliefs is

$$E(u(a_i,\omega)) = \sum_{\omega\in\Omega} u(a_i,\omega)\mu(\omega) = \sum_{\omega_i\in X} \omega_i \mu_i(\omega_i)$$

Independence implies that this expression is independent of any other act. The same is true of the expected value of the normalized utilities $z(a,\omega) \equiv \exp(u(a,\omega)/\lambda)$. We use Ez^{a_i} to refer to this expectation, so

$$Ez^a = \sum_{\omega_i\in X} \exp(\omega_i/\lambda)\mu_i(\omega_i)$$

Let $P(a)$ denote the probability that an action a is chosen under the optimal scheme and let $B = \{a \in A | P^a > 0\}$ denote the set of chosen acts. Suppose that there exist two actions b and c such that $Ez^b < Ez^c$ and yet b is chosen and c is not chosen. Define

$$\Sigma(\omega) = \left(\sum_{a\in B} P(a)z(a,\omega)\right)^{-1}$$

The necessary and sufficient conditions imply

$$Ez^b\Sigma = \sum_{\omega\in\Omega}\left[\frac{z(b,\omega)}{\Sigma(\omega)}\mu(\omega)\right] = 1 \geq \sum_{\omega\in\Omega}\left[\frac{z(c,\omega)}{\Sigma(\omega)}\mu(\omega)\right] = Ez^c\Sigma$$

The fact that c is not chosen means that $c \notin B$, which implies that Σ_ω is independent of $z(c,\omega)$, as $\Sigma(\omega)$ is a function of $z(a,\omega)$ for every $a \in B$, all of which are independent of $z(c,\omega)$. It follows that

$$Ez^c\Sigma = Ez^cE\Sigma$$

Furthermore, by assumption

$$Ez^cE\Sigma > Ez^bE\Sigma$$

Implying

$$Ez^b\Sigma > Ez^bE\Sigma$$

which can only hold if

$$\mathrm{cov}(z^b,\Sigma) > 0$$

Now

$$
\begin{aligned}
\operatorname{cov}(z^b,\Sigma) &\equiv E\frac{z^b - Ez^b}{\sum_{a\in B} P(a)z(a,\omega)} \\
&= E\frac{z^b - Ez^b}{P(b)z(b,\omega)+\sum_{a\in B/b} P(a)z(a,\omega)} \\
&= E\left\{E\left\{\frac{z^b - Ez^b}{P(b)z(b,\omega)+\sum_{a\in B/b} P(a)z(a,\omega)}\middle| \sum_{a\in B\setminus b} P(a)z(a,\omega)\right\}\right\} \\
&= E\left\{\operatorname{cov}\left(z^b, \frac{1}{P(b)z(b,\omega)+\sum_{a\in B/b} P(a)z(a,\omega)}\right)\middle| \sum_{a\in B\setminus b} P(a)z(a,\omega)\right\}
\end{aligned}
$$

where where the second to last line follows form the law of iterated expectations and the last line follows form the independence of b and the other elements of B which implies that Ez^b is independent of $\sum_{a\in B\setminus b} P(a)z(a,\omega)$. For given $\sum_{a\in B\setminus b} P(a)z(a,\omega)$ it is clear that $\operatorname{cov}\left(z^b, \frac{1}{P(b)z(b,\omega)+\sum_{a\in B/b} P(a)z(a,\omega)}\right)$ is negative. As the expectation of a collection of negative numbers is negative

$$\operatorname{cov}(z^b,\Sigma) < 0$$

This contradiction establishes the proof.

Let (μ, A) be an independent consumption problem and $\{a_1, ...a_N\} = B \subset A$ be a set of ex ante identical actions (i.e. $\mu_i(x) = \mu_j(x) = \mu_B(x)$ for all $x \in X$ and $i,j \leq N$). Then a strategy that picks each $a_i \in B$ with the unconditional probability $\frac{1}{N}$ and assigns unconditional probabilities according to equation 2 is optimal if, for each $a_j \notin B$

$$\sum_{x\in X} \exp(x/\lambda)\mu_j(x) \leq \frac{1}{n}\left[\sum_{\bar{x}\in X^N} \frac{\Pi_{n=1}^{N}\mu_B(\bar{x}_n)}{\sum_{n=1}^{N} \exp\left(\bar{x}_n/\lambda\right)}\right]^{-1}$$

Proof of Lemma 1 First we show that a strategy of choosing $a \in B$ with probability $\frac{1}{N}$ satisfies the necessary and sufficient conditions of Proposition 1 for all $a \in B$. Given the proposed strategy and the independence of the decisions

$$\sum_{\omega\in\Omega} \frac{z(a,\omega)\mu(\omega)}{\sum_{b\in B} P(b)z(b,\omega)} = \sum_{\omega\in\Omega} \frac{z(a,\omega)\Pi_{i=1}^{M}\mu_i(\omega_i)}{\sum_{b\in B} \frac{1}{N}z(b,\omega)}$$

for all $a \in B$. The denominator only depends only on the first N elements of the state vector, so we can sum across the other states giving

$$\sum_{x \in X^N} \frac{\exp\left(x_a/\lambda\right) \Pi_{i=1}^N \mu_B(x_i)}{\frac{1}{N}\left[\sum_{i=1}^N \exp\left(x_i/\lambda\right)\right]} \tag{10}$$

We now subdivide X^N into equivalence classes as follows: given $x_1, x_2 \in X^N$ then x_1 and x_2 are members of the same equivalence class $\hat{X}$ if x_2 is a permutation of x_1. We rewrite (10)

$$\sum_{\hat{X}} \sum_{x \in \hat{X}} \frac{\exp\left(x_a/\lambda\right) \Pi_{i=1}^N \mu_B(x_i)}{\frac{1}{N}\left[\sum_{i=1}^N \exp\left(x_i/\lambda\right)\right]}$$

Note that conditional on $x \in \hat{X}$, the denominator, $\sum_{i=1}^N \exp\left(x_i/\lambda\right)$, and the probability, $\Pi_{i=1}^N \mu_B(x_i)$, are constant. Therefore

$$\sum_{\hat{X}} \sum_{x \in \hat{X}} \frac{\exp\left(x_a/\lambda\right) \Pi_{i=1}^N \mu_B(x_i)}{\frac{1}{N}\left[\sum_{i=1}^N \exp\left(x_i/\lambda\right)\right]} = \sum_{\hat{X}} \frac{\Pi_{i=1}^N \mu_B(x_i) \sum_{x \in \hat{X}} \exp\left(x_a/\lambda\right)}{\frac{1}{N}\left[\sum_{i=1}^N \exp\left(x_i/\lambda\right)\right]}$$

$$\begin{aligned} & \sum_{\hat{X}} \sum_{x \in \hat{X}} \frac{N \exp\left(x_a/\lambda\right) \mu(\hat{X})}{|\hat{X}| \left[\sum_{i=1}^N \exp\left(x_i/\lambda\right)\right]} \\ = \quad & \sum_{\hat{X}} \mu(\hat{X}) \frac{N \sum_{x \in \hat{X}} \exp\left(x_a/\lambda\right)}{|\hat{X}| \sum_{i=1}^N \exp\left(x_i/\lambda\right)} \end{aligned}$$

Let $\mu(\hat{X})$ denote the probability $x \in \hat{X}$. We have $\mu(\hat{X}) = |\hat{X}| \Pi_{i=1}^N \mu_B(x_i)$. Note also that $\sum_{x \in \hat{X}} \exp\left(x_a/\lambda\right)$ and $\sum_{i=1}^N \exp\left(x_i/\lambda\right)$ share common terms in common proportions. The only difference is the frequency in which these terms appear. $\sum_{x \in \hat{X}} \exp\left(x_a/\lambda\right)$ has $|\hat{X}|$ terms and $\sum_{i=1}^N \exp\left(x_i/\lambda\right)$ has N terms. $\sum_{x \in \hat{X}} \exp\left(x_a/\lambda\right) / \sum_{i=1}^N \exp\left(x_i/\lambda\right)$ is therefore equal to $|\hat{X}|/N$. Substituting these results, we have

$$\sum_{\hat{X}} \frac{\Pi_{i=1}^N \mu_B(x_i) \sum_{x \in \hat{X}} \exp\left(x_a/\lambda\right)}{\frac{1}{N}\left[\sum_{i=1}^N \exp\left(x_i/\lambda\right)\right]} = \sum_{\hat{X}} \frac{\frac{\mu(\hat{X})}{|\hat{X}|} |\hat{X}|}{\frac{1}{N} N} = \sum_{\hat{X}} \mu(\hat{X}) = 1$$

So that the conditions of Proposition 1 hold for $a \in B$.

To complete the proof we need to show that for all $a \in A \backslash B$

$$\sum_{\omega \in \Omega} \frac{z(a,\omega)\mu(\omega)}{\sum_{b \in A} P(b) z(b,\omega)} \leq 1$$

Given independence and the proposed strategy this becomes

$$\sum_{x \in X} \exp(x/\lambda)\mu_a(x) \sum_{\bar{x} \in X^N} \frac{\Pi_{i=1}^N \mu_B(\bar{x}_i)}{\frac{1}{N}\sum_{i=1}^N \exp\left(\bar{x}_i/\lambda\right)} \leq 1$$

This relationship holds given the conditions of the lemma. This completes the proof of the Lemma.