

Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model

By FILIP MATĚJKA AND ALISDAIR MCKAY*

Individuals must often choose among discrete actions with imperfect information about their payoffs. Before choosing, they have an opportunity to study the payoffs, but doing so is costly. This creates new choices such as the number of and types of questions to ask. We model these situations using the rational inattention approach to information frictions. We find that the decision maker's optimal strategy results in choosing probabilistically in line with a generalized multinomial logit model, which depends both on the actions' true payoffs as well as on prior beliefs. Keywords: discrete choice, information, rational inattention, multinomial logit.

Economists and psychologists have long known that scarce attention plays an important role in decision making (Simon, 1959; Kahneman, 1973). In this paper we study the discrete choice behavior of an agent who must allocate his limited attention to the available information about the choice situation.

It is not uncommon for one to be faced with a choice among discrete actions with imperfect information about the payoffs of each. Before making a choice, one may have an opportunity to study the actions and their payoffs, however, in most cases it is too costly to investigate to the point where the payoffs are known with certainty. As a result, some uncertainty about the payoffs remains when one chooses among the actions even if complete information were available in principle. Because of this uncertainty, the action that is ultimately chosen may not be the one that provides the highest payoff to the decision maker. Moreover, the noise in the decision process may lead identical individuals to make different choices. In this manner, imperfect information naturally leads choices to contain errors and be probabilistic as opposed to deterministic.

In this context, the decision maker faces choices of how much to study the environment and what to investigate when doing so. That is, the decision maker must choose how to allocate his attention. For example, a firm might choose how long to spend interviewing candidates for a job and choose what to ask them during the interview. After completing the interview, the firm faces a discrete choice among the candidates.

We explore the optimal “information processing” behavior of a decision maker for whom acquiring information is costly and characterize the resulting choice behavior in this discrete choice context. As choices are probabilistic, our characterization involves describing the probability with which the decision maker selects a particular action in a particular choice problem. Specifically, we model the cost of acquiring and processing information using the rational inattention framework introduced by Sims (1998, 2003) which uses information theory to measure the amount of information the decision maker processes.

* Matějka: CERGE-EI, A joint workplace of the Center for Economic Research and Graduate Education, Charles University, and the Economics Institute of the Academy of Sciences of the Czech Republic. Politických vězňů 7, Praha 1, 150 00, the Czech Republic, filip.matejka@cerge-ei.cz. McKay: Department of Economics, Boston University, 270 Bay State Road, Boston, Massachusetts, USA, amckay@bu.edu. We thank Michal Bauer, Levent Celik, Satyajit Chatterjee, Faruk Gul, Christian Hellwig, Stepan Jurajda, Bart Lipman, Tony Marley, Andreas Ortmann, Juan Ortner, Christopher Sims, Leonidas Spiliopoulos, Jakub Steiner, Jorgen Weibull, and Michael Woodford for comments and helpful discussions. This research was funded by GAČR P402/11/P236, by European Research Council under the European Community's Seventh Framework Programme FP7/2007-2013 grant agreement N. 263790, and by a grant RRC 11+001 from the CERGE-EI Foundation under a program of the Global Development Network. All opinions expressed are those of the authors and have not been endorsed by CERGE-EI or the GDN.

The major appeal of the rational inattention approach to information frictions is that it does not impose any particular assumptions on what agents learn or how they go about learning it. Instead, the rational inattention approach derives the information structure from the utility-maximizing behavior of the agents for whom information is costly to acquire. As a result, rationally inattentive agents process information that they find useful and ignore information that is not worth the effort of acquiring and processing.

Our main finding is that the decision maker’s optimal information processing strategy results in probabilistic choices that follow a logit model that reflects both the actions’ true payoffs as well as the decision maker’s prior beliefs. In a choice among N actions with any form of prior beliefs, our modified logit formula takes the form

$$(1) \quad \frac{e^{(v_i+\alpha_i)/\lambda}}{\sum_{j=1}^N e^{(v_j+\alpha_j)/\lambda}},$$

where v_i is the payoff of action i and λ is a parameter that scales the cost of information. The decision maker’s prior knowledge and information processing strategy are incorporated into the choice probabilities through the weights, α_i , attached to each action. These action weights shift the choice probabilities towards those actions that appeared to be good candidates *a priori* and they are completely independent of the actual payoffs of the actions. They do, however, depend on the cost of information, λ . As the cost of information rises, the action weights are revised so that the decision maker’s choice becomes less sensitive to the actual payoffs of the actions and more sensitive to his prior beliefs.

When the *a priori* beliefs do not influence the choice, our work provides a new foundation for the multinomial logit model. We show that whenever the actions are exchangeable in the decision maker’s prior, α_i is constant across i so equation (1) simplifies to the standard multinomial logit formula. Cases where the actions are homogenous *a priori* arise naturally whenever the decision maker lacks specific knowledge that allows him to distinguish between the actions before entering the choice situation.

The multinomial logit model is perhaps the most commonly used model of discrete choice and it has two canonical foundations.² According to the random utility derivation, the decision maker evaluates the payoffs of the actions with some noise. If the noise in the evaluation is additively separable and independently distributed according to the extreme value distribution, then the multinomial logit model emerges.³ To date, there is not a clear economic or psychological justification of why these disturbances should be extreme value distributed.

The multinomial logit model has a second canonical foundation, namely, Luce’s (1959) derivation from the independence of irrelevant alternatives (IIA) axiom, which states that ratios of choice probabilities are independent of the choice set. In Section III we show that the modified logit formula in equation (1) can also be derived from two axioms that are related to the IIA axiom, but modified in a way that reflects the role of prior beliefs in the rational inattention model.

We believe our findings are important for two reasons. First, while the logit model is sometimes used in situations in which information frictions are thought to be an important part of the choice environment, there has not previously been a fully-specified model of those information frictions that justifies the use of the multinomial logit. We fill that gap and demonstrate that the fully-specified

²The logit model was first proposed for binary choices by Bradley and Terry (1952) and the multinomial logit was introduced by Luce (1959). Anderson, de Palma and Thisse (1992), McFadden (2001), and Train (2009) present surveys of discrete choice theory and the multinomial logit model.

³Luce and Suppes (1965, p. 338) attribute this result to Holman and Marley (unpublished). See McFadden (1974) and Yellott (1977) for the proof that a random utility model generates the logit model only if the noise terms are extreme value distributed.

model is not subject to some of the criticisms of the logit model. Second, most existing work with rational inattention has focussed on situations where the decision maker chooses from a continuous set of actions. In this context, the model remains tractable if one assumes the agent is acquiring information about a normally-distributed quantity and the objective function is quadratic, as under these assumptions the decision maker chooses normally distributed signals. Beyond this situation, however, the continuous-choice rational inattention model must be solved numerically and even numerical solutions can be difficult to obtain. In contrast, we show here that the discrete-choice version of the rational inattention model is extremely tractable. Our results allow the rational inattention framework to be easily applied by building on a large body applied theoretical work that exploits the tractability of the multinomial logit.⁴

The rationally inattentive agent’s choices are context dependent where context is defined by prior beliefs that represent the immediately apparent characteristics of the actions. As the decision maker’s prior changes, his choice behavior changes both due to standard Bayesian updating and through endogenous changes in his information processing strategy. Changes in the information processing strategy can lead to results that are appealingly intuitive and to results that seem surprising. We show:

- The rationally inattentive agent ignores duplicate actions, which are known *a priori* to yield the same payoff.
- Adding an action to the choice set can increase the likelihood that an existing action is selected—an outcome that cannot occur in any random utility model.
- Action i is more likely to be selected if the decision maker’s beliefs about it improve in the sense of first-order-stochastic dominance. As we demonstrate, this monotonicity does not always hold under Bayesian updating with an exogenous information structure.

That the rationally inattentive agent ignores duplicate actions stands in contrast to the standard logit model. Debreu (1960) criticized IIA and the logit model for having counter-intuitive implications for the treatment of duplicate actions. The standard model treats each action as distinct and only allows them to be differentiated in one dimension (e.g. their payoffs) and as a result there is no sense of similarity between actions. Adding a duplicate of one action will therefore increase the probability that the action (or its duplicate) is selected. In our setting, however, we can introduce the concept of similarity through prior knowledge that they have the same payoff even if this common payoff is unknown. We show that when a duplicate action is added to the choice set, the rationally inattentive agent will choose to ignore it. Therefore the model does not display the counter-intuitive behavior that Debreu (1960) criticized.

The paper is organized as follows: In the remainder of this section we review related work. Section I presents the choice setting and discusses the assumptions underlying the rational inattention approach to information frictions. Section II studies the decision maker’s optimal strategy. Section III presents a characterization result that connects the optimal behavior of the rationally inattentive agent to versions of Luce’s IIA axiom. Section IV demonstrates how the decision maker’s prior knowledge influences his choice behavior. Finally, Section V concludes.

⁴The multinomial logit model is commonly used in the industrial organization and international trade literatures as a model of consumer demand, in political economy models of voting, and in experimental economics to capture an element of bounded rationality in subject behavior. See Anderson, de Palma and Thisse (1992) for a survey of its use in industrial organization. The logit demand structure was introduced to international trade by Goldberg (1995) and Verboven (1996). The logit model was incorporated into probabilistic voting models by Lindbeck and Weibull (1987). Work following McKelvey and Palfrey (1995) uses logit response functions to capture randomness in the responses of experimental subjects playing a game. Matějka and McKay (2012) is an example of how the results of this paper can be integrated with existing results based on the multinomial logit model to study the optimal price-setting behavior of firms facing rationally inattentive consumers.

RELATED LITERATURE

Our work relates to the literature on rational inattention. Most existing work with rational inattention has focussed on situations where the decision maker chooses from a continuous choice set. Rational inattention has mostly been applied in macroeconomic contexts such as consumption-savings problems (Sims, 2006; Maćkowiak and Wiederholt, 2010) and price setting (Mackowiak and Wiederholt, 2009; Matějka, 2010a).⁵ A few papers, however, consider applications with binary choice problems. Woodford (2009) was the first to do so in a study of a binary choice of whether to adjust a price, while Yang (2011) investigates a global game setting with the choice of whether to invest or not. Moreover, Matějka and Sims (2010) and Matějka (2010a) provide a connection between the continuous and discrete problems by showing that rationally inattentive agents can voluntarily constrain themselves to a discrete choice set even when the initial set of available actions is continuous. We extend the existing literature by establishing a connection between rational inattention and the multinomial logit model and characterize the implications of rational inattention for discrete choice behavior.⁶

Weibull, Mattsson and Voorneveld (2007) study a discrete choice model with imperfect information in which the decision maker receives signals centered on the true payoffs of the actions and is able to control the precision, but not the shape of the signal. The authors show that when the signals are i.i.d. extreme-value distributed the best policy is to select the action with the highest signal. It then follows from the usual random utility derivation that the choice probabilities follow the logit model.⁷ Our assumptions differ in three ways: first, our decision maker is not restricted to a particular distribution for signals and has complete freedom in selecting the nature of information he receives. Second, we allow the actions to differ in the decision maker's prior beliefs leading to the incorporation of prior knowledge into the choice probabilities. Third, our information processing cost has different qualitative properties that, for instance, imply higher payoffs never lead to lower expected utility, which is not the case in Weibull, Mattsson and Voorneveld. In their setup, better payoffs can make some states more difficult to infer and can lower the expected utility the decision maker achieves.

Closely related to our work is that of Caplin and Dean (2013a) who provide an alternative method for characterizing solutions in a related environment. Masatlioglu, Nakajima and Ozbay (2012) and Manzini and Mariotti (2013) model imperfect attention using a consideration set approach. Under this approach, choices occur in two stages: first, some of the available actions are selected into a consideration set and then the utility maximizing action is chosen from the consideration set. Under this approach, the decision maker may overlook some of the available actions while in our setting the decision maker is aware of all actions, but may not be aware of their exact characteristics.

The rational inattention approach to information frictions uses information theoretic concepts to measure the amount of information processed by the decision maker and there is a mathematical connection between the entropy function, which is at the heart of information theory, and the multinomial logit. This connection has appeared in the context of statistical estimation (Anas, 1983) and in the context of an agent stabilizing a trembling hand (Stahl, 1990; Mattsson and Weibull, 2002). The mathematical connection between entropy and the logit or Maxwell-Boltzmann distribution was shown by Jaynes (1957) who reinterpreted statistical mechanics in physics using information theory.⁸ What is new in our work is that we are considering the decision problem of an agent who

⁵Other applications are Luo (2008); Luo and Young (2009); Tutino (2009); Van Nieuwerburgh and Veldkamp (2010); Mondria (2010); Matějka (2010b); Paciello and Wiederholt (2014); Stevens (2011).

⁶In an independent paper that is as of yet unfinished, Woodford (2008) notices the connection to the logit model in the context of a binary choice problem, but does not explore the connection in further detail.

⁷Natenzon (2010) studies a related model with correlated Gaussian signals.

⁸Another derivation of the logit distribution is due to Shannon (1959) who derived the multinomial logit formula in an engineering application that is the dual to our problem in the *a priori* homogeneous case. Shannon's question is how quickly

must acquire information about the payoffs of the available actions. The context of information-constrained choice allows us to investigate how prior beliefs affect the choice probabilities and what conditions on them are needed for the standard logit model to arise. The decision-making interpretation and the presence of prior beliefs result in somewhat different optimization problems from those in physics and engineering applications.

I. The model

We first present the formal model and then discuss its interpretation in Section I.B.

A. Framework

The decision maker chooses an action from the set $A = \{1, \dots, N\}$. The state of nature is a vector $\mathbf{v} \in \mathbb{R}^N$ where v_i is the payoff of action $i \in A$. The decision maker has imperfect information about the state of nature and so is unsure of the payoff that results from each action. The decision maker observes a signal on the state and then chooses an action as a Bayesian expected utility maximizer. Before receiving the signal, the decision maker is able to choose an information processing strategy that determines the joint distribution of the signal and the state. Processing information is costly and more informative strategies are more costly.

The decision problem has two stages. In the first stage, the decision maker selects an information strategy to refine his belief about the state. The second stage is a standard choice under uncertainty with the beliefs generated in the first stage.

In the second stage, the decision maker has a belief $B \in \Delta(\mathbb{R}^N)$, where $\Delta(\mathbb{R}^N)$ is the set of all probability distributions on \mathbb{R}^N . Given this belief, the decision maker chooses the action with the highest expected payoff

$$(2) \quad V(B) \equiv \max_{i \in A} \mathbb{E}_B[v_i],$$

where $V : \Delta(\mathbb{R}^N) \rightarrow \mathbb{R}$. Let $a : \Delta(\mathbb{R}^N) \rightarrow A$ be the action strategy that satisfies $a(B) = \arg \max_{i \in A} \mathbb{E}_B[v_i]$.⁹

In the first stage of the decision problem, the decision maker selects an information strategy. The decision maker is endowed with some prior belief $G \in \Delta(\mathbb{R}^N)$ about the state and can receive signals $\mathbf{s} \in \mathbb{R}^N$ on the state to update his beliefs. The information strategy is a joint distribution $F(\mathbf{s}, \mathbf{v}) \in \Delta(\mathbb{R}^{2N})$ of signals and states such that the marginal distribution over states equals the decision maker's prior G , which ensures that the decision maker's posterior beliefs are consistent with his prior. Given this restriction, the decision maker is only free to select the conditional distribution $F(\mathbf{s}|\mathbf{v})$. The other conditional distribution, $F(\mathbf{v}|\mathbf{s})$, is the posterior belief after receiving signal \mathbf{s} .¹⁰

In the first stage, the decision maker maximizes the *ex ante* expected payoff from the second stage less the cost of information, $\hat{c}(F)$, which reflects the cost of generating signals of different

a message can be communicated through a limited-capacity channel, such as a telegraph wire, without distorting the message beyond a certain degree on average. We thank Michael Woodford for pointing us to this connection to Shannon's work and an anonymous referee for pointing us to that of Jaynes.

⁹If there are multiple maximizers we let $a(B)$ denote an arbitrary element of $\arg \max_{i \in A} \mathbb{E}_B[v_i]$.

¹⁰A signal about the state is normally modeled as a variable with a distribution that varies with the state. A family of such conditional distributions along with the prior over the state implies a joint distribution of the signal and the state. It is more convenient to work directly with this joint distribution, $F(\mathbf{s}, \mathbf{v}) \in \Delta(\mathbb{R}^{2N})$, as the information strategy.

degrees of precision. The decision maker solves

$$(3) \quad \max_{F \in \Delta(\mathbb{R}^{2N})} \int_{\mathbf{v}} \int_{\mathbf{s}} V(F(\cdot|\mathbf{s})) F(d\mathbf{s}|\mathbf{v}) G(d\mathbf{v}) - \hat{c}(F),$$

$$(4) \quad \text{such that} \quad \int_{\mathbf{s}} F(d\mathbf{s}, \mathbf{v}) = G(\mathbf{v}) \quad \forall \mathbf{v} \in \mathbb{R}^N,$$

where V is determined by (2).

We assume the entropy-based cost function used in the rational inattention literature

$$(5) \quad \hat{c}(F) \equiv \lambda \left(H(G) - E_{\mathbf{s}}[H(F(\cdot|\mathbf{s}))] \right),$$

where the parameter $\lambda \geq 0$ is the unit cost of information and $H(B)$ denotes the uncertainty of belief B as measured by its entropy. Our model allows for either a continuous or discrete distribution of states. For the discrete case, entropy is

$$H(B) = - \sum_k P_k \log(P_k),$$

where P_k is the probability of state k .¹¹ Entropy is a measure of uncertainty that measures the average unlikeliness of events. For a discrete distribution with M equally-probable events, the entropy is $\log(M)$. When M increases, each event becomes less likely and entropy increases. Similarly, the entropy of a Gaussian distribution is an increasing function of its variance. $\hat{c}(F)$ is therefore proportional to the expected difference between the prior uncertainty about the state and the posterior uncertainty after observing the signal. This difference is referred to as the mutual information between the state and the signal. We discuss the motivation for this cost function in Section I.B.

A choice problem is summarized by the prior beliefs about the state, G , and the cost of information processing, λ . In the first stage, the decision maker selects the information strategy F knowing the expected payoffs, V , and action strategy, a , that will prevail in the second stage.

DEFINITION 1: (Model: state - signal - action). *Let $G(\mathbf{v})$ be the decision maker's prior and let λ be the unit cost of information. The discrete choice strategy of the rationally inattentive decision maker is the pair (F, a) of the information and action strategies such that a is given by the maximization in (2) and F solves (3).*

We now show that the state-contingent choice behavior of the rationally inattentive decision maker can be found as the solution to a simpler maximization problem that does not make reference to signals or posterior beliefs.

The combination of an information strategy and an action strategy together induce a joint distribution between the selected action and the state. For any strategy (F, a) , let $S_i \equiv \{\mathbf{s} \in \mathbb{R}^N : a(F(\mathbf{v}|\mathbf{s})) = i\}$ be the set of signals that lead to the action i . Let $\mathcal{P}_i(\mathbf{v})$ be the induced probability of selecting i conditional on state \mathbf{v} , which is given by:

$$(6) \quad \mathcal{P}_i(\mathbf{v}) \equiv \int_{\mathbf{s} \in S_i} F(d\mathbf{s}|\mathbf{v}).$$

¹¹For a distribution that has a pdf f entropy is $-\int_{\mathbf{v}} f(\mathbf{v}) \log f(\mathbf{v}) d\mathbf{v}$.

Let \mathcal{P} denote the collection $\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ and let

$$(7) \quad \mathcal{P}_i^0 \equiv \int_{\mathbf{v}} \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v})$$

be the unconditional probability of selecting action i .

An important property of rationally inattentive behavior is that under an optimal strategy each action is selected in at most one posterior. Receiving distinct signals that lead to the same action is inefficient as information is acquired but not acted upon. This behavior follows from the convexity of the entropy-based cost function.¹²

LEMMA 1: *If (F, a) is an optimal strategy, then $\forall i \in A$ such that $\mathcal{P}_i^0 > 0$, there exists a posterior $B_i \in \Delta(\mathbb{R}^N)$ such that conditional on taking action i , the probability that the posterior belief is B_i equals 1. That is, $F(\mathbf{v}|\mathbf{s}) = B_i$ for all $\mathbf{s} \in S_i$ up to a set of measure zero under F . Moreover, the expected payoff in (3) and the cost of information (5) can be calculated from \mathcal{P} as:*

$$(8) \quad \int_{\mathbf{v}} \int_{\mathbf{s}} V(F(\cdot|\mathbf{s})) F(d\mathbf{s}|\mathbf{v}) G(d\mathbf{v}) = \sum_{i=1}^N \int_{\mathbf{v}} v_i \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}),$$

$$(9) \quad \hat{c}(F) = c(\mathcal{P}, G) \equiv \lambda \left(- \sum_{i=1}^N \mathcal{P}_i^0 \log \mathcal{P}_i^0 + \int_{\mathbf{v}} \left(\sum_{i=1}^N \mathcal{P}_i(\mathbf{v}) \log \mathcal{P}_i(\mathbf{v}) \right) G(d\mathbf{v}) \right).$$

PROOF:

See Appendix A.

As each action is associated with a particular signal, the cost of information can be calculated as the mutual information between states and actions as shown in equation (9). We can now write the model in terms of state-contingent choice probabilities alone.

COROLLARY 1: (**Model: state - action**). *The collection of conditional probabilities $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ is induced by a solution to the problem in Definition 1 if and only if it solves the following optimization problem.*

$$(10) \quad \max_{\mathcal{P}=\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N} \sum_{i=1}^N \int_{\mathbf{v}} v_i \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}) - c(\mathcal{P}, G),$$

subject to

$$(11) \quad \forall i : \quad \mathcal{P}_i(\mathbf{v}) \geq 0 \quad \forall \mathbf{v} \in \mathbb{R}^N,$$

$$(12) \quad \sum_{i=1}^N \mathcal{P}_i(\mathbf{v}) = 1 \quad \forall \mathbf{v} \in \mathbb{R}^N,$$

and where $c(\mathcal{P}, G)$ is defined by (9).

PROOF:

See Appendix A.

In what follows we study solutions to the problem in Corollary 1 knowing that a solution to the problem in Definition 1 can be constructed from it.¹³

¹²Convexity is a natural property of the cost function. For instance, it is implied by any cost function for which an information structure that is more informative in the sense of Blackwell (1953) is more costly.

¹³To construct such a solution, let $\{s_i\}_{i=1}^N$ be N distinct points in \mathbb{R}^N , i.e. signals; let $F(s_i, \mathbf{v})$ be the product of $\mathcal{P}_i(\mathbf{v})$ and $G(\mathbf{v})$; and let $a(F(\mathbf{v}|s_i)) \equiv i$. Note that the solution to Definition 1 is not unique as the supports of the signals, $\{s_i\}_{i=1}^N$, can be selected arbitrarily, but the induced \mathcal{P} is given by the solution to Corollary 1. If the solution to the problem in Corollary 1 is unique, then any solution to Definition 1 will induce the same joint distribution of states and actions, \mathcal{P} .

B. Interpretation

In our model, the decision maker is rationally inattentive meaning he chooses: i) how much attention to pay, i.e. how much information to process; ii) what to pay attention to, i.e. what information to process; and iii) what action to select given the new information. The benefit of processing information is that it allows the decision maker to more accurately select actions with high payoffs and the details of the choice problem determine the optimal information strategy.

To fix ideas, let us consider a simple example in which the decision maker must choose between taking a red bus, a blue bus, or a train. The decision maker knows how long the journey takes by train and this results in a payoff of R for sure. He is less familiar with the buses and does not know whether they are fast or slow. Suppose a fast bus results in the payoff of 1 and a slow bus yields 0. The decision maker's prior is that either bus is fast with probability $1/2$ and their speeds have a correlation of ρ . The decision maker can acquire more information about the buses by checking the schedule, asking others for their experiences, and so on. The choice problem is summarized in Table 1.

	<u>state 1</u>	<u>state 2</u>	<u>state 3</u>	<u>state 4</u>
action 1: red bus	0	1	0	1
action 2: blue bus	0	0	1	1
action 3: train	R	R	R	R
$G(\mathbf{v})$	$\frac{1}{4}(1 + \rho)$	$\frac{1}{4}(1 - \rho)$	$\frac{1}{4}(1 - \rho)$	$\frac{1}{4}(1 + \rho)$

TABLE 1—PAYOFF MATRIX, EACH COLUMN LISTS \mathbf{v} FOR THAT STATE.

Depending on the parameters ρ and R , the decision maker will choose to process more or less information in total and will choose to receive signals that are more or less informative about particular states of nature. For example if $R \geq 1$, the decision maker need not process any information and can simply select the train. If $\rho = -1$ and $R = 1/2$, the decision maker knows that the train is always inferior to one of the buses and will devote his attention to determining which bus is faster. If $\rho = 1$, the decision maker knows the buses yield the same payoff and so he can ignore one of them. We discuss this example in more detail in Section IV.

The decision maker's information strategy can generate any distribution of signals conditional on each state. For instance one candidate strategy in the example above is a signal that takes two values and perfectly distinguishes states 1 and 3 from states 2 and 4. Such a signal would inform the decision maker of the payoff of the red bus, but not that of the blue bus. This (sub-optimal) strategy requires a one bit reduction of entropy when $\rho = 0$. Another one-bit strategy could distinguish between all four states, but only imperfectly meaning there is noise in the signal. Which strategy is optimal depends on the payoffs of the actions in each of the states and the distribution of noise in the optimal strategy will be related to the costs of the mistakes the noise leads to.

At the heart of the model is a formulation of the cost of asking questions, querying a database or otherwise obtaining and processing information. The decision maker is completely free in choosing what questions to ask. If the decision maker so chooses he can learn very precise details about some aspect of the environment and completely ignore other aspects. The only consideration is that any reduction in uncertainty is costly. In the terms of the model, the choice of the conditional distribution $F(\mathbf{s}|\mathbf{v})$ is unconstrained.

Our results depend crucially on the cost function we use, but this cost function arises naturally in a particular context as we now explain. Suppose the decision maker acquires information by asking a series of yes-or-no questions and suppose there is a cost per question he asks. The more questions the decision maker asks, the tighter his posterior beliefs can be. The entropy-based cost function is

proportional to the expected number of questions needed to implement the information strategy.¹⁴ Alternatively, it is possible to motivate entropy as a value of information based on willingness to pay for an information structure as in Cabrales, Gossner and Serrano (2013) or as a limit on what can be communicated between players of a game as in Gossner, Hernandez and Neyman (2006).

While it is important that information processing is measured with the expected entropy reduction, it is not crucial that the total cost is proportional to this quantity. What is important for the behavior of choice probabilities across states is only the marginal cost of a unit of information at the optimum. If we were to allow for a more general cost of information, the results in Section II would only change by replacing λ with the marginal cost of information, but otherwise would remain unchanged. A similar argument applies to the specification used in Sims (2003), where the decision maker faces a fixed information capacity or in other words a cost function that jumps from zero to infinity at a particular level of information processing. In that case, λ would be replaced by the shadow cost of information. These different formulations are not equivalent when we consider changes of behavior across different choice problems as in Section IV where the decision maker might change the amount of information he processes. The formulation we use, in contrast to the fixed information capacity formulation, implies that the decision maker can choose the amount of information to process and, for instance, pay more attention when the stakes are high.

II. Solving the model

A. Solving for choice probabilities

We begin our analysis of the model with a general analytical expression for the probability that the decision maker chooses a particular action conditional on the state of nature.

THEOREM 1: *If $\lambda > 0$, then the optimal strategy has conditional choice probabilities that satisfy:*

$$(13) \quad \mathcal{P}_i(\mathbf{v}) = \frac{\mathcal{P}_i^0 e^{v_i/\lambda}}{\sum_{j=1}^N \mathcal{P}_j^0 e^{v_j/\lambda}} \quad \text{almost surely.}$$

If $\lambda = 0$, then the decision maker selects the action(s) with the highest payoff with probability one.

PROOF:

See Appendix B.

(13) holds almost surely as opposed to point-wise because the decision maker's objective is unaffected by deviations in his strategy on a measure-zero set of states. We can understand several properties of the decision maker's behavior from equation (13). The unconditional probabilities, \mathcal{P}_i^0 , are by definition independent of the state. They are the marginal probabilities of selecting each action before the agent starts processing any information and they only depend on the prior knowledge $G(\mathbf{v})$ and the cost of information λ .

When the unconditional probabilities are uniform, $\mathcal{P}_i^0 = 1/N$ for all i , (13) becomes the usual multinomial logit formula. As we discuss in Section II.B, this happens when G is invariant to permutations of its arguments. In other cases, the conditional choice probabilities are not driven just by \mathbf{v} , as in the logit case, but also by the unconditional probabilities of selecting each action, $\{\mathcal{P}_i^0\}_{i=1}^N$. These unconditional probabilities are monotonic transformations of the action weights that we referred to in the introduction with $\exp(\alpha_i/\lambda) = \mathcal{P}_i^0$. Using this transformation, equation (13) can be rewritten as equation (1) and the choice probabilities can be interpreted as a multinomial logit in which the payoff of action i is shifted by the term α_i . When an action seems very attractive

¹⁴See the coding theorem of information theory (Shannon, 1948; Cover and Thomas, 2006).

a priori, then it has a relatively high probability of being selected even if its true payoff is low. As the cost of information, λ , rises, the less the decision maker finds out about the state of nature and the more he decides based on prior knowledge of the actions. In what follows, we find it more convenient to work directly with the unconditional probabilities, \mathcal{P}_i^0 , rather than the transformed version, α_i .

The parameter λ converts bits of information to utils. Therefore, if one scales the payoffs of all of the actions by a constant c , while keeping the information cost, λ , fixed, the problem is equivalent to the one with the original payoffs and the information cost scaled by $1/c$. By scaling up the payoffs, one is scaling up the differences between them and therefore raising the stakes for the decision maker. The decision maker chooses to process more information because more is at stake and thus is more likely to select the action that provides the highest payoff. The decision maker behaves just as he would if the cost of information had fallen.

Equation (13) does not give a fully explicit expression for the choice probabilities because it depends on the \mathcal{P}_i^0 terms, which are themselves part of the decision maker's strategy although they are independent of the state. We can substitute equation (13) into the objective function to arrive at the following formulation of the optimization problem that we will use to solve for the unconditional choice probabilities.

LEMMA 2: *Alternative formulation:* *The collection of conditional probabilities $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ solves (10)-(12) in Corollary 1 if and only if the probabilities are given by (13) with the unconditional probabilities, $\{\mathcal{P}_i^0\}_{i=1}^N$, that solve:*

$$(14) \quad \max_{\{\mathcal{P}_i^0\}_{i=1}^N} \int_{\mathbf{v}} \lambda \log \left(\sum_{i=1}^N \mathcal{P}_i^0 e^{v_i/\lambda} \right) G(d\mathbf{v}).$$

subject to

$$(15) \quad \forall i : \quad \mathcal{P}_i^0 \geq 0,$$

$$(16) \quad \sum_i \mathcal{P}_i^0 = 1.$$

PROOF:

See Appendix B.

This novel formulation is useful for two reasons. First, it allows for clearer insights into the properties of choice probabilities than the original problem and we use it extensively in the proofs in Section IV. Second, it greatly reduces the complexity of the optimization problem and allows for more efficient computations. Rational inattention problems with continuous choice variables can also be formulated this way.¹⁵ The first-order conditions of this problem give us

COROLLARY 2: *Normalization condition:* *For all i such that $\mathcal{P}_i^0 > 0$, the solution satisfies*

$$(17) \quad \int_{\mathbf{v}} \frac{e^{v_i/\lambda}}{\sum_{j=1}^N \mathcal{P}_j^0 e^{v_j/\lambda}} G(d\mathbf{v}) = 1.$$

PROOF:

¹⁵For a general problem under rational inattention, such as in Sims (2006), the alternative formulation takes the same form with v_i replaced by $U(X, Y)$, with the sum over i replaced by an integral over Y and \mathbf{v} replaced by X , where X is the unknown and Y the action.

See online Appendix C.

We call this the normalization condition because if one multiplies both sides of equation (17) by \mathcal{P}_i^0 , the result ensures that the conditional choice probabilities in equation (13) integrate to the unconditional choice probability as in (7). The analysis in the next subsection and in Section IV is based on finding solutions to equations (13) and (17).

Online Appendix C establishes that a solution to the maximization problem in Corollary 1 exists, however, the solution may not be unique. Cases with multiple solutions require a special structure for the uncertainty in which the payoffs co-move across states in a very rigid way. For instance, when the payoffs of two actions are equal in all states of the worlds, then the decision maker can relocate the choice probabilities between the two actions arbitrarily and realize the same expected payoff. Perhaps an illustrative interpretation of non-unique solutions is that when there are multiple solutions there always exists at least one action that can be eliminated from the choice set without reducing the expected utility that the decision maker can achieve. These eliminations can be repeated until the solution is unique. Online Appendix C provides the exact conditions that are necessary and sufficient for uniqueness.

Theorem 1 provides a solution for the choice probabilities given a set of unconditional choice probabilities, $\{\mathcal{P}_i^0\}_{i=1}^N$. There is no general closed-form solution for the unconditional probabilities, but in some particular cases we can establish several properties of the solution. First, in Section II.B we show that when the decision maker does not have information to distinguish between the actions *a priori* then $\mathcal{P}_i^0 = 1/N$ for all i . Second, if one action is dominated by another in all states of the world, then it will never be selected.¹⁶ Third, for the extreme case $\lambda = 0$ the prior is irrelevant and the decision maker selects the highest-payoff action in all states of the world. Conversely, if $\lambda = \infty$, the actual realization of the state is irrelevant and the decision maker selects the action with the highest expected payoff according to the prior. The decision maker may also behave in this way for a finite λ if he *chooses* not to process any information. In addition to these properties, Section IV establishes some comparative static properties that describe how the unconditional probabilities change as prior beliefs change. Finally, Lemma 2 offers a convenient means of computing the unconditional probabilities in an application.

B. Multinomial logit

We now present conditions under which the behavior of the rationally inattentive agent follows the multinomial logit model. This connection to the logit model holds across different realizations of the state.

Let us assume that all the actions are exchangeable in the prior G , i.e. the prior is invariant to all permutations of the entries of \mathbf{v} . We call such actions *a priori homogeneous*. The actions will be *a priori* homogeneous whenever the decision maker does not distinguish between them before he starts processing information, which is a plausible benchmark case.

PROBLEM 1: *The decision maker chooses $i \in \{1, \dots, N\}$, where the actions are a priori homogeneous and take different payoffs with positive probability.*

PROPOSITION 1: **Logit:** *In Problem 1, the probability of choosing action i as a function of the state \mathbf{v} is*

$$(18) \quad \mathcal{P}_i(\mathbf{v}) = \frac{e^{v_i/\lambda}}{\sum_{j=1}^N e^{v_j/\lambda}}.$$

¹⁶This is evident from the optimization problem in Lemma 2 where it is clear that if $v_i > v_j$ for all states of the world then the objective function is always increased by reducing \mathcal{P}_j^0 and relocating that probability to \mathcal{P}_i^0 .

PROOF:

See Appendix B.

This is the multinomial logit formula written in terms of exponentials as it is most often used in practice. We show that the *a priori* homogeneity of the actions implies that the unconditional probabilities are uniform so that (13) then takes the form of the logit as the only thing that distinguishes actions is their actual payoffs. The assumption on the possibility of different payoffs is needed for uniqueness.

Let us emphasize that $\mathcal{P}_i(\mathbf{v})$ does not depend on the prior G in this case. As long as the actions are *a priori* homogeneous, the resulting choice probabilities take the form of (18). This feature is particularly useful as it makes applications of the rational inattention framework very simple in this case. This result follows from the endogenous information structure in the model. The optimal choice of information fixes the nature of the optimal posteriors and the decision maker selects what information to acquire so as to arrive at posteriors of that form. In contrast, with an exogenous information structure, changes in the prior lead directly to changes in the posterior, with corresponding effects on the choice probabilities.

III. Characterization and IIA

Suppose an analyst has data on choice probabilities as a function of the state of the world. How can she judge whether they are consistent with the modified multinomial logit formula in equation (13) without knowing the payoffs of the actions or the decision maker’s prior beliefs? A specific example helps fix ideas: suppose the decision maker has a choice among three actions, the consequences of which depend on the state of nature, and the consequences take the form of a particular good that the decision maker receives as listed in Table 2. If the decision maker chooses action 1 in state 1 he receives a mug and so on. Suppose the analyst observes the consequences of the actions and the choice probabilities of the actions across these different states of the world.¹⁷

	state 1	state 2	state 3
action 1	mug	pen	nothing
action 2	mug	pen	pen
action 3	nothing	t-shirt	t-shirt

TABLE 2—STATE-CONTINGENT CONSEQUENCES.

In the rational inattention model, the implication of Theorem 1 is that while the decision maker processes information about the unknown state, the choice probabilities are as if the decision maker knows the state, attaches weights $\{\mathcal{P}_i^0\}_{i=1}^N$ to the actions, and then chooses among the actions probabilistically according to equation (13). In this section we establish that two conditions on choice probabilities are necessary and sufficient for the existence of weights $\{\mathcal{P}_i^0\}_{i=1}^N$ and payoffs for each consequence such that the choice probabilities can be represented through equation (13). This alternative derivation of equation (13) is closely related to the Luce (1959) derivation of the standard multinomial logit model from the IIA axiom, which is an alternative derivation to the random utility derivation.

Let us first introduce the standard IIA axiom. Let X be a set of consequences. The IIA axiom states that for any two sets $A \subseteq X$ and $B \subseteq X$ that both contain elements x and y , and if

¹⁷This type of data on state-contingent choice probabilities can be collected in a laboratory setting as Caplin and Dean (2013b) have done.

$\Pr(y \text{ chosen from A}) > 0$, then

$$(19) \quad \frac{\Pr(x \text{ chosen from A})}{\Pr(y \text{ chosen from A})} = \frac{\Pr(x \text{ chosen from B})}{\Pr(y \text{ chosen from B})}.$$

Luce (1959) showed that if the choice probabilities satisfy this axiom¹⁸ then there is a function $v : X \rightarrow \mathbb{R}$ such that the choice probabilities can be expressed through the multinomial logit formula

$$\Pr(x \text{ chosen from A}) = \frac{e^{v(x)}}{\sum_{y \in A} e^{v(y)}}.$$

Luce viewed IIA as a natural benchmark for choice behavior. Following Luce’s work, Debreu (1960) criticized IIA for generating counter-intuitive implications. We show in the next section that this criticism does not apply to the rational inattention model.

Our axioms differ from IIA by allowing for the possibility that the probability of selecting a consequence can depend on the action that leads to the consequence. This is motivated by the rational inattention model where the decision maker selects actions depending on the realized state, but also on based on his prior knowledge, and thus prior beliefs affect choice probabilities in addition to the consequences of the actions.

Let $\mathbf{x} \in X^N$ be a vector of consequences where the i ’th entry, x_i , is the consequence that the decision maker receives if he selects action i . In the spirit of Section I, we will call \mathbf{x} the state. Slightly abusing notation, let $\mathcal{P}_i(\mathbf{x})$ be the probability of selecting action i in state \mathbf{x} . Suppose we observe choice probabilities from state \mathbf{x} and from another state \mathbf{y} such that actions i and j yield the same consequences in the two states: $x_i = y_i$ and $x_j = y_j$. Just like IIA, Axiom 1 states that the ratio of choice probabilities for actions i and j is constant across these states.

AXIOM 1: Independence from irrelevant alternatives (consequences): *if $\mathcal{P}_j(\mathbf{x}) > 0$, then*

$$(20) \quad \frac{\mathcal{P}_i(\mathbf{x})}{\mathcal{P}_j(\mathbf{x})} = \frac{\mathcal{P}_i(\mathbf{y})}{\mathcal{P}_j(\mathbf{y})} \quad \forall \mathbf{x}, \mathbf{y}, i, j \text{ s.t. } x_i = y_i \text{ and } x_j = y_j.$$

However, Axiom 1 is weaker than the standard IIA axiom (19), since it requires the equality between the ratio of probabilities only when the two consequences in the two states follow from the same pair of actions. In the example in Table 2, Axiom 1 requires that $\frac{\mathcal{P}_2(\text{state 2})}{\mathcal{P}_3(\text{state 2})} = \frac{\mathcal{P}_2(\text{state 3})}{\mathcal{P}_3(\text{state 3})}$, the ratio of probabilities of actions 2 and 3 does not change if only the consequence from action 1 changes. The standard IIA axiom would imply more restrictions. For example, $\frac{\mathcal{P}_1(\text{state 2})}{\mathcal{P}_3(\text{state 2})}$ would equal the ratios above as well.

Our second axiom is a restriction on how choice probabilities vary across actions independently of the associated consequences. If two actions yield the same consequence, i.e. $x_i = x_j$, then the ratio $\mathcal{P}_i(\mathbf{x})/\mathcal{P}_j(\mathbf{x})$ is independent of the particular consequence x_i and independent of the consequences of other actions $k \notin \{i, j\}$.

AXIOM 2: Independence from irrelevant alternatives (actions): *if $\mathcal{P}_j(\mathbf{x}) > 0$, then*

$$(21) \quad \frac{\mathcal{P}_i(\mathbf{x})}{\mathcal{P}_j(\mathbf{x})} = \frac{\mathcal{P}_i(\mathbf{y})}{\mathcal{P}_j(\mathbf{y})} \quad \forall \mathbf{x}, \mathbf{y}, i, j \text{ s.t. } x_i = x_j \text{ and } y_i = y_j.$$

In the example, Axiom 2 requires that $\frac{\mathcal{P}_1(\text{state 1})}{\mathcal{P}_2(\text{state 1})} = \frac{\mathcal{P}_1(\text{state 2})}{\mathcal{P}_2(\text{state 2})}$, the ratio of probabilities of two actions does not change if they both lead to the same consequence in each state (state 1: mug,

¹⁸Luce considered an axiom that is equivalent to IIA as we have presented it.

state 2: pen). This ratio is independent of what the consequence is (i.e. mug or pen) and is independent of the consequence of action 3. In Axiom 2 the ratio of choice probabilities might differ from unity even though both actions lead to the same consequence.

We will now show that there is a weight for each action and a payoff for each of the consequences such that the choice probabilities can be written as in equation (13) if and only if the choice probabilities satisfy Axioms 1 and 2. We first establish the following:

LEMMA 3: *If Axioms 1 and 2 hold for a choice among $N \geq 3$ actions, then each action is either never selected for any \mathbf{x} or it is selected with positive probability for all \mathbf{x} .*

PROOF:

See online Appendix D.

We will refer to actions that are always selected with some probability as “positive actions.” The main result of this section is:

PROPOSITION 2: *Let there exist at least three positive actions. Then the choice probabilities satisfy Axioms 1 and 2 if and only if there exist non-negative constants $\{\mathcal{P}_i^0\}_{i=1}^N$ such that $\sum_{i=1}^N \mathcal{P}_i^0 = 1$ and also a function $v : X \rightarrow \mathbb{R}$ such that for any i and any $\mathbf{x} \in X^N$, the probability of selecting action i , is*

$$(22) \quad \mathcal{P}_i(\mathbf{x}) = \frac{\mathcal{P}_i^0 e^{v(x_i)}}{\sum_j \mathcal{P}_j^0 e^{v(x_j)}}.$$

PROOF:

See online Appendix D.

To relate equation (22) to equation (13) we set $v_i = \lambda v(x_i)$. The ratio in equation (21) is then equal to $\mathcal{P}_i^0 / \mathcal{P}_j^0$ from the rational inattention model and the ratio in equation (20) is equal to $\mathcal{P}_i^0 e^{v_i/\lambda} / [\mathcal{P}_j^0 e^{v_j/\lambda}]$.

IV. The role of prior beliefs

In this section, we present comparative static results that demonstrate how the choice probabilities that solve the decision problem in Definition 1 depend on the decision maker’s prior beliefs about the payoffs. Mathematically, we show how changes in the prior distribution $G(\mathbf{v})$ affect the unconditional probabilities or in other words the action weights.

We begin with general results concerning two main types of modifications to the prior: changes in the level of payoffs and changes in the co-movement of payoffs across states. First, we show that anytime an action is improved in the prior then the decision maker is more likely to select it. We then show that if two actions become more similar, defined as having payoffs that co-move more strongly across states, then the probability that the decision maker selects either of these actions falls. We connect this notion of similarity to Debreu’s criticisms of the IIA property of the multinomial logit model and show that the rationally inattentive agent’s behavior is not subject to the same criticisms. Finally, we demonstrate that the behavior of the rationally inattentive agent fails regularity as adding an additional action to the choice set can increase the probability that an existing action is selected. An implication of this failure of regularity is that the rationally inattentive agent cannot generally be viewed as maximizing a random utility function.

A. Monotonicity

The following proposition states that a change in the prior that makes one action more attractive *ceteris paribus* leads the decision maker to select that action with a higher unconditional probability

and therefore a higher probability in all states of the world. Such a result does not generally hold under Bayesian updating if the information structure is exogenously given.¹⁹ In this way, rational inattention places more structure on choice behavior than does incompleteness of information alone.

PROPOSITION 3: *Assume $\lambda > 0$ and let $\{\mathcal{P}_i^0\}_{i=1}^N$ be the unique solution to the agent’s maximization problem with prior $G(\mathbf{v})$. If $\hat{G}(\hat{\mathbf{v}})$ is generated from $G(\mathbf{v})$ by transforming \mathbf{v} to $\hat{\mathbf{v}}$ such that $\hat{v}_i = v_i$ for all $i > 1$, $\hat{v}_1 \geq v_1$ for all \mathbf{v} and $\hat{v}_1 > v_1$ on a set of positive measure, then $\hat{\mathcal{P}}_1^0 \geq \mathcal{P}_1^0$, where $\{\hat{\mathcal{P}}_i^0\}_{i=1}^N$ is the solution to the problem with prior \hat{G} . The last inequality holds strictly if $\mathcal{P}_1^0 \in (0, 1)$.*

PROOF:

See online Appendix E.

That the monotonicity is strict only when $\mathcal{P}_1^0 \in (0, 1)$ is intuitive. When $\mathcal{P}_1^0 = 1$, there is no scope for increasing it further. When $\mathcal{P}_1^0 = 0$, action 1 might be so unattractive to start with that the improvement does not lead the decision maker to select it with any probability.

B. Similarity and independence from irrelevant alternatives

While the previous proposition considers the effect of improving one action, suppose we leave the marginal distributions of the payoffs of each action unchanged, but change the prior to increase the degree to which the payoffs of two actions co-move. What happens to the choice probabilities in this case? To motivate this investigation, consider Debreu’s famous criticism of IIA, equation (19), which is now known as the red-bus blue-bus problem. The well known example goes: The decision maker is pairwise indifferent between choosing a bus or a train, and selects each with probability 1/2. If a second bus of a different color is added to the choice set and the decision maker is indifferent to the color of the bus, then IIA—and therefore the multinomial logit—implies probabilities for the train and the two buses of 1/3, 1/3, 1/3. Debreu argued that this is counter-intuitive. This is counter-intuitive because the two buses are the same consequence and adding the second bus only provides the decision maker with another way of arriving at the same consequence and payoff as an existing action. It would be more intuitive if adding the second bus did not affect the probability of choosing the train, for example resulting in choice probabilities of 1/2, 1/4, 1/4.

We will say that two actions are duplicates if the prior asserts that the payoffs of the two actions are equal to one another in all states of the world: it is clear to the decision maker *a priori* that it is irrelevant which of the two actions he selects. We will now show that the rationally inattentive agent does not display the counter-intuitive behavior that Debreu criticized and then expand these ideas to cases in which the actions can be thought to be “similar” i.e. have payoffs that are correlated across states of the world, but are not exact duplicates.

¹⁹ Here is a counterexample in which improving prior beliefs about one action leads it to be selected less often under Bayesian updating with an exogenous information structure. Imagine that there are two states of the world that each occur with probability 1/2. There are two available actions with the following payoffs in the two states of the world

	state 1	state 2
v_1	-2	4
v_2	0	x .

Suppose the information structure is that the decision maker receives a signal $y = v_2 + \varepsilon$ where the noise ε is distributed uniformly on $(-1, 1)$. When $x = -2$, the signal perfectly distinguishes between the states of the world and so the decision maker selects action 1 in state 2 and action 2 in state 1. When action 2 is improved so that $x = 0$, the signal contains no information about the state of the world. In this case, the decision maker always selects action 1 as it has the larger expected payoff. So improving action 2 leads it to be selected less often.

DUPLICATES

We study a generalized version of Debreu’s bus problem to analyze how the rationally inattentive agent treats duplicate actions. Duplicates carry the same payoff although this common payoff may be unknown.

DEFINITION 2: *Actions i and j are duplicates if and only if the probability that $v_i \neq v_j$ is zero.*

PROBLEM 2: *The decision maker chooses $i \in \{1, \dots, N + 1\}$, where the actions N and $N + 1$ are duplicates.*

The following proposition states that duplicate actions are treated as a single action. We compare the choice probabilities in two choice problems, where the second one is constructed from the first by duplicating one action. In the first problem, the decision maker’s prior is $G(\mathbf{v})$, where $\mathbf{v} \in \mathbb{R}^N$. In the second problem, the decision maker’s prior is $\hat{G}(\mathbf{u})$, where $\mathbf{u} \in \mathbb{R}^{N+1}$. \hat{G} is generated from G by duplicating action N . This means that actions N and $N + 1$ satisfy Definition 2, and $G(\mathbf{v})$ is the marginal of $\hat{G}(\mathbf{u})$ with respect to u_{N+1} .

PROPOSITION 4: *If $\{\mathcal{P}_i^0\}_{i=1}^N$ and $\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ are unconditional and conditional choice probabilities that are a solution to the choice problem with prior G , then $\{\hat{\mathcal{P}}_i(\mathbf{u})\}_{i=1}^{N+1}$ solve the corresponding problem with the added duplicate of action N if and only if they satisfy the following:*

$$(23) \quad \hat{\mathcal{P}}_i(\mathbf{u}) = \mathcal{P}_i(\mathbf{v}), \quad \forall i < N$$

$$(24) \quad \hat{\mathcal{P}}_N(\mathbf{u}) + \hat{\mathcal{P}}_{N+1}(\mathbf{u}) = \mathcal{P}_N(\mathbf{v}),$$

where $\mathbf{v} \in \mathbb{R}^N$ and $\mathbf{u} \in \mathbb{R}^{N+1}$, and $v_k = u_k$ for all $k \leq N$. The analogous equalities hold for the unconditional probabilities.

PROOF:

See online Appendix E.

The implication of this proposition is that the decision maker treats duplicate actions as though they were a single action. The behavior of the rationally inattentive agent does not always satisfy IIA and as a result is not subject to Debreu’s critique. As we showed in Section III, a version of IIA holds for a fixed prior, but not when the prior changes.

SIMILAR ACTIONS

The case of exact duplicates is somewhat extreme as the decision maker knows *a priori* that the payoffs of the two actions are exactly equal. Here we consider actions with payoffs that are correlated, but not identical. We show that the probability that the decision maker selects either of two actions (among three or more) decreases as those two actions become more similar. By “more similar” we mean that we consider an alternative choice situation in which we increase the probability that the two actions have the same payoff by shifting the prior probability from states of the world in which their payoffs differ to states of the world where their payoffs are the same. We have the following proposition.

PROPOSITION 5: *Assume $\lambda > 0$ and let $\{\mathcal{P}_i^0\}_{i=1}^N$ be the unique solution to the agent’s maximization problem with prior $G(\mathbf{v})$. Let the prior $\hat{G}(\cdot)$ be such that there exist two states of positive probabilities P_1 and P_2 with $(v_1, v_2) = (H, L)$ in state 1, and $(v_1, v_2) = (L, H)$ in state 2, and let the payoff of any other action be equal in both states.*

If $\hat{G}(\hat{\mathbf{v}})$ is generated from $G(\cdot)$ by relocating probability mass $\Pi \leq \min(P_1, P_2)$ from state 1 to state 3, where $(v_1, v_2) = (L, L)$, and relocating probability mass Π from state 2 to state 4, where $(v_1, v_2) = (H, H)$, then $\hat{\mathcal{P}}_1^0 + \hat{\mathcal{P}}_2^0 < \mathcal{P}_1^0 + \mathcal{P}_2^0$, where $\{\hat{\mathcal{P}}_i^0\}_{i=1}^N$ is the solution to the problem with prior \hat{G} .

PROOF:

See online Appendix E.

Intuitively, a higher degree of co-movement in the manner described in the proposition means that the event $(v_1 = H) \cup (v_2 = H)$ has lower probability and this is the event in which the decision maker is most interested in selecting action 1 or 2. In the next section, we present an example of this type of effect in which increasing the correlation between the payoffs of two actions decreases their cumulative probability.

C. Examples

We bring our analysis to a close with two examples. First, we demonstrate the effect of similarity among actions. We then show that adding an additional action can increase the probability that an existing action is selected.

CO-MOVEMENT

We now return to the example from Section I.B where the decision maker chooses between a red bus, a blue bus and a train. We set the payoff of the train, R , to $1/2$ so all three actions have the same *a priori* expected payoff. Nevertheless, for some values of ρ the decision maker will ignore one of the actions completely. This example demonstrates that when the allocation of attention is endogenous, the decision maker can choose to investigate different actions in different levels of detail.

PROBLEM 3: *The decision maker chooses from the set {red bus, blue bus, train} with payoffs given in Table 1.*

In online Appendix F we describe how to solve the problem analytically. Figure 1 illustrates the behavior of the model for various values of ρ and λ . The figure shows the unconditional probability that the decision maker selects a bus of a given color (the probability is the same for both buses). As the correlation between the payoffs of the buses decreases, the probability that a bus carries the largest payoff among the three actions increases and the unconditional probability of choosing either bus increases, too. If the buses' payoffs are perfectly correlated, then the sum of their probabilities is 0.5, they are effectively treated as one action, i.e. they become duplicates in the limit. On the other hand, if $\rho = -1$, then the unconditional probability of either bus is 0.5 and thus the train is never selected.

For $\lambda > 0$ and $\rho \in (-1, 1)$, the probability that a bus is selected is larger than it is in the perfect information case ($\lambda = 0$). With a larger cost of information, the decision maker economizes on information by paying more attention to choosing among the buses and less to assessing their payoffs relative to the train.

The choice probabilities strongly reflect the endogeneity of the information structure in this case. As the correlation decreases, the decision maker knows that the best action is more likely to be one of the buses. As a result, the decision maker focusses more of his attention on choosing between the buses and eventually ignores the train completely. Notice that this can happen even when there is some chance that the train actually yields a higher payoff.

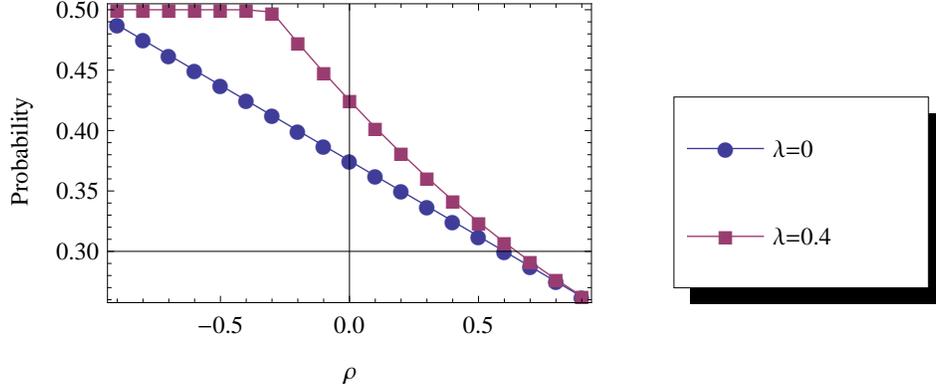


FIGURE 1. UNCONDITIONAL PROBABILITY OF SELECTING A BUS FOR VARIOUS VALUES OF λ AND ρ . THE PROBABILITY IS THE SAME FOR BOTH THE RED AND BLUE BUSES.

FAILURE OF REGULARITY

Random utility models, such as the standard multinomial logit model, have the feature that adding an additional action to the choice set will not increase the probability that an existing action is selected (Luce and Suppes, 1965, p. 342). However, the following example demonstrates that the behavior of the rationally inattentive agent does not always satisfy this regularity condition.

PROBLEM 4: *Suppose there are three actions and two states of nature. The actions have the following payoffs in the two states:*

	<i>state 1</i>	<i>state 2</i>
<i>action 1</i>	0	1
<i>action 2</i>	1/2	1/2
<i>action 3</i>	Y	$-Y$
<i>prior prob.</i>	$g(1)$	$g(2) = 1 - g(1)$

First, consider a variant of this choice problem in which only actions 1 and 2 are available. In online Appendix F we show that there exists $g(1) \in (0, 1)$ large enough that the decision maker will not process information and will select action 2 with probability one in all states of nature so $\mathcal{P}_1^0 = 0$. Now add action 3 to the choice set. For a large enough value of Y and the given $g(1) \in (0, 1)$, the decision maker will find it worthwhile to process information about the state in order to determine whether action 3 should be selected. Given that the decision maker will now have information about the state, if state 2 is realized, the decision maker might as well select action 1. From an *a priori* perspective, there is a positive probability of selecting action 1 so $\mathcal{P}_1^0 > 0$. The choice probabilities conditional on the realization of the state are given by equation (13), which implies that the probability of selecting action 1 is zero if $\mathcal{P}_1^0 = 0$ and positive if $\mathcal{P}_1^0 > 0$ and all actions have finite payoffs. So we have the following.

PROPOSITION 6: *For $\lambda > 0$, there exist $g(1) \in (0, 1)$ and $Y > 0$ such that adding action 3 to the choice set in Problem 4 increases the probability that action 1 is selected in all states of nature.*

PROOF:

See online Appendix F.

COROLLARY 3: *The behavior of a rationally inattentive agent cannot always be described by a random utility model.*

Obviously there are cases, such as the standard logit case, when the rationally inattentive agent's behavior *can* be described by a random utility model.

V. Conclusion

This paper builds on the literature that treats scarce attention as an important factor in choice behavior. Rational inattention is an appealing model of attention allocation because it does not depend on assumptions about how the decision maker allocates his attention to the available information except in the form of a well-founded cost function that measures the decision maker's information processing effort. While appealing, the rational inattention model in its original continuous-action form is considered by some to be intractable. In this paper, however, we have shown that the discrete choice behavior of a rationally inattentive agent has a simple analytical structure providing a new foundation for the multinomial logit model. The standard logit model emerges in cases where the actions are homogeneous *a priori*. More generally, choices depend on the context formed by prior beliefs. While the decision maker's prior can be a complicated infinite-dimensional object, its effect on choice probabilities is captured by a simple vector of action weights. The resulting tractability allows us to establish results on monotonicity, co-movement, and uniqueness, which are general and fundamental properties of rationally inattentive behavior. The discrete choice framework presented here facilitates the application of rational inattention to new questions.

REFERENCES

- Anas, Alex.** 1983. “Discrete choice theory, information theory and the multinomial logit and gravity models.” *Transportation Research Part B: Methodological*, 17(1): 13–23.
- Anderson, Simon P., Andre de Palma, and Jacques-Francois Thisse.** 1992. *Discrete Choice Theory of Product Differentiation*. Cambridge, MA:MIT Press.
- Blackwell, D.** 1953. “Equivalent comparison of experiments.” *Annals of Mathematical Statistics*, , (24).
- Bradley, Ralph Allan, and Milton E. Terry.** 1952. “Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons.” *Biometrika*, 39(3/4): pp. 324–345.
- Cabrales, Antonio, Olivier Gossner, and Roberto Serrano.** 2013. “Entropy and the Value of Information for Investors.” *The American Economic Review*, 103(1): 360–377.
- Caplin, Andrew, and Mark Dean.** 2013*a*. “The Behavioral Implications of Rational Inattention with Shannon Entropy.”
- Caplin, Andrew, and Mark Dean.** 2013*b*. “Rational Inattention and State Dependent Stochastic Choice.” New York University working paper.
- Cover, Thomas M., and Joy A. Thomas.** 2006. *Elements of Information Theory*. Hoboken, NJ:Wiley.
- Debreu, Gerard.** 1960. “Review of Individual Choice Behavior by R. D. Luce.” *American Economic Review*, 50(1).
- Goldberg, Pinelopi Koujianou.** 1995. “Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry.” *Econometrica*, 63(4).
- Gossner, Olivier, Penelope Hernandez, and Abraham Neyman.** 2006. “Optimal use of communication resources.” *Econometrica*, 74(6): 1603–1636.
- Jaynes, Edwin T.** 1957. “Information theory and statistical mechanics.” *Physical review*, 106(4): 620.
- Kahneman, Daniel.** 1973. *Attention and Effort*. New Jersey:Prentice Hall.
- Lindbeck, Assar, and Jrgen W. Weibull.** 1987. “Balanced-budget redistribution as the outcome of political competition.” *Public Choice*, 52.
- Luce, R. Duncan.** 1959. *Individual Choice Behavior: a Theoretical Analysis*. New York:Wiley.
- Luce, R. Duncan, and Patrick Suppes.** 1965. “Preference, Utility, and Subjective Probability.” In *Handbook of Mathematical Psychology*. Vol. 3, , ed. R. Luce, R. D.; Bush and E. Galanter, 249–410. New York:Wiley.
- Luo, Yulei.** 2008. “Consumption dynamics under information processing constraints.” *Review of Economic Dynamics*, 11(2).
- Luo, Yulei, and Eric R. Young.** 2009. “Rational Inattention and Aggregate Fluctuations.” *The B.E. Journal of Macroeconomics*, 9(1).
- Maćkowiak, Bartosz Adam, and Mirko Wiederholt.** 2010. “Business Cycle Dynamics under Rational Inattention.” C.E.P.R. Discussion Papers CEPR Discussion Papers 7691.

- Mackowiak, Bartosz, and Mirko Wiederholt.** 2009. "Optimal Sticky Prices under Rational Inattention." *The American Economic Review*, 99.
- Manzini, Paola, and Marco Mariotti.** 2013. "Stochastic Choice and Consideration Sets." Forthcoming in *Econometrica*.
- Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Y. Ozbay.** 2012. "Revealed Attention." *The American Economic Review*, 102(5): 2183–2205.
- Matějka, Filip.** 2010a. "Rationally Inattentive Seller: Sales and Discrete Pricing." CERGE-EI Working Papers wp408.
- Matějka, Filip.** 2010b. "Rigid Pricing and Rationally Inattentive Consumer." CERGE-EI Working Papers wp409.
- Matějka, Filip, and Alisdair McKay.** 2012. "Simple Market Equilibria with Rationally Inattentive Consumers." *American Economic Review*, 102(3): 24 – 29.
- Matějka, Filip, and Christopher A. Sims.** 2010. "Discrete Actions in Information-Constrained Tracking Problems." Princeton University Technical report.
- Mattsson, Lars-Goran, and Jorgen W. Weibull.** 2002. "Probabilistic Choice and Procedurally Bounded Rationality." *Games and Economic Behavior*, 41(1).
- McFadden, Daniel.** 1974. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, ed. Paul Zarembka. New York:Academic Press.
- McFadden, Daniel.** 2001. "Economic Choices." *The American Economic Review*, 91(3): pp. 351–378.
- McKelvey, Richard D., and Thomas R. Palfrey.** 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, 10(1).
- Mondria, Jordi.** 2010. "Portfolio choice, attention allocation, and price comovement." *Journal of Economic Theory*, 145(5).
- Natenzon, Paulo.** 2010. "Random Choice and Learning." Princeton University Working Paper.
- Paciello, Luigi, and Mirko Wiederholt.** 2014. "Exogenous Information, Endogenous Information, and Optimal Monetary Policy." *The Review of Economic Studies*, 81(1): 356–388.
- Shannon, C. E.** 1959. "Coding Theorems for a Discrete Source With a Fidelity Criterion." Vol. 7.
- Shannon, Claude E.** 1948. "A Mathematical Theory of Communication." *The Bell System Technical Journal*, 27.
- Simon, Herbert A.** 1959. "Theories of Decision-Making in Economics and Behavioral Science." *The American Economic Review*, 49(3): pp. 253–283.
- Sims, Christopher A.** 1998. "Stickiness." *Carnegie-Rochester Conference Series on Public Policy*, 49.
- Sims, Christopher A.** 2003. "Implications of Rational Inattention." *Journal of Monetary Economics*, 50(3).
- Sims, Christopher A.** 2006. "Rational Inattention: Beyond the Linear-Quadratic Case." *The American Economic Review*, 96(2).

- Stahl, D. O.** 1990. "Entropy Control Costs and Entropic Equilibria." *International Journal of Game Theory*, 19(2).
- Stevens, Luminita.** 2011. "Price Adjustment in a Model with Multiple Price Policies." *Working paper, Columbia University*.
- Train, Kenneth.** 2009. *Discrete Choice Methods with Simulation*. Cambridge, U.K.:Cambridge University Press.
- Tutino, Antonella.** 2009. "The rigidity of choice: Lifetime savings under information-processing constraints." Working paper.
- Van Nieuwerburgh, Stijn, and Laura Veldkamp.** 2010. "Information Acquisition and Under-Diversification." *Review of Economic Studies*, 77(2).
- Verboven, Frank.** 1996. "International Price Discrimination in the European Car Market." *RAND Journal of Economics*, 27(2).
- Weibull, Jorgen W., Lars-Goran Mattsson, and Mark Voorneveld.** 2007. "Better May be Worse: Some Monotonicity Results and Paradoxes in Discrete Choice Under Uncertainty." *Theory and Decision*, 63.
- Woodford, Michael.** 2008. "Inattention as a Source of Randomized Discrete Adjustment." Columbia University working paper.
- Woodford, Michael.** 2009. "Information-constrained state-dependent pricing." *Journal of Monetary Economics*, 56(Supplement 1).
- Yang, Ming.** 2011. "Coordination with Rational Inattention." Princeton University Working Paper.
- Yellott, John I.** 1977. "The relationship between Luce's Choice Axiom, Thurstone's Theory of Comparative Judgment, and the double exponential distribution." *Journal of Mathematical Psychology*, 15(2).

Proof of Lemma 1. By contradiction. Let us assume that there exists an action i such that $\mathcal{P}_i^0 > 0$ for which there are distinct posteriors realized with a positive probability. Formally, let S_i be the set of signals \mathbf{s} that lead to action i and suppose there exist $S_i^1, S_i^2 \subseteq S_i$ such that $\int_{\mathbf{v}} \int_{\mathbf{s} \in S_i^k} F(d\mathbf{v}, d\mathbf{s}) > 0$ $k = 1, 2$ and $\forall \mathbf{s}_1 \in S_i^1, \forall \mathbf{s}_2 \in S_i^2 : F(\mathbf{v}|\mathbf{s}_1) \neq F(\mathbf{v}|\mathbf{s}_2)$. We can construct another feasible strategy that generates the same expected payoff at a lower information processing cost than F . The new strategy, $\bar{F} \in \Delta(\mathbb{R}^N)$, is generated from the original strategy by relocating the probability mass from $S_i^1 \cup S_i^2$ to a single signal, call it $\hat{\mathbf{s}} \in S_i^1 \cup S_i^2$. We only alter the conditional distribution $F(\mathbf{s}|\mathbf{v})$ and do not change the marginal over \mathbf{v} , which implies the constraint (4) is still satisfied. Intuitively, we are scrambling the signal so that the decision maker does not observe events in $S_i^1 \cup S_i^2$ separately, but just observes that one of them has occurred.

By assumption, action i is optimal conditional on observing $\mathbf{s} \in S_i^1$ and also for $\mathbf{s} \in S_i^2$. By the law of iterated expectations it follows that it is optimal to select action i conditional on observing $\hat{\mathbf{s}}$ or equivalently $\mathbf{s} \in S_i^1 \cup S_i^2$:

$$\mathbb{E}[v_i|\hat{\mathbf{s}}] = \mathbb{E}\left[\mathbb{E}\left[v_i|\mathbf{s} \in S_i^k\right]|\hat{\mathbf{s}}\right] \geq \mathbb{E}\left[\mathbb{E}\left[v_j|\mathbf{s} \in S_i^k\right]|\hat{\mathbf{s}}\right] = \mathbb{E}[v_j|\hat{\mathbf{s}}] \quad \forall j \in A.$$

Therefore the new action strategy satisfies $\bar{a}(\bar{F}(\cdot|\hat{\mathbf{s}})) = i$. For any other signal $\mathbf{s} \neq \hat{\mathbf{s}}$, the conditional distribution $\bar{F}(\cdot|\mathbf{s})$ remains equal to $F(\cdot|\mathbf{s})$ and so $\bar{a}(\bar{F}(\cdot|\mathbf{s})) = a(F(\cdot|\mathbf{s}))$ for all $\mathbf{s} \neq \hat{\mathbf{s}}$.

The constructed strategy (\bar{F}, \bar{a}) generates the same expected payoff as the original strategy. The expected payoff can be written as

$$\int_{\mathbf{v}} \int_{\mathbf{s} \notin S_i^1 \cup S_i^2} V(F(\cdot|\mathbf{s}))F(d\mathbf{s}|\mathbf{v})G(d\mathbf{v}) + \int_{\mathbf{v}} \int_{\mathbf{s} \in S_i^1 \cup S_i^2} V(F(\cdot|\mathbf{s}))F(d\mathbf{s}|\mathbf{v})G(d\mathbf{v}).$$

The first term is unaffected by the change of strategy as neither conditional distribution $F(\mathbf{v}|\mathbf{s})$ nor $F(\mathbf{s}|\mathbf{v})$ is altered on this set of signals. Under both strategies, the second term can be expressed as the product of $\Pr(\mathbf{s} \in S_i^1 \cup S_i^2)$ and the expected payoff of action i conditional on $\mathbf{s} \in S_i^1 \cup S_i^2$, which follows from the law of iterated expectations as $V(F(\cdot|\mathbf{s})) = \mathbb{E}[v_i|\mathbf{s}]$ on this set. Neither the probability of the set nor the conditional expectation given the set is affected by the change of strategy.

As entropy is a strictly concave function of the distribution, the cost $\hat{c}(\bar{F})$ is lower than $\hat{c}(F)$,²⁰ and thus (F, a) cannot be a solution since (\bar{F}, \bar{a}) generates the same expected payoff at a lower information cost. This proves that posteriors leading to the same action cannot differ with positive probability.

To show (8), we use the implication from above that $F(\mathbf{v}|\mathbf{s})$ is constant for all $\mathbf{s} \in S_i$ and the fact that $V(F(\cdot|\mathbf{s}))$ is a function of \mathbf{s} :

$$\begin{aligned} \int_{\mathbf{v}} \int_{\mathbf{s}} V(F(\cdot|\mathbf{s}))F(d\mathbf{s}|\mathbf{v})G(d\mathbf{v}) &= \sum_{i=1}^N V(F(\cdot|S_i)) \int_{\mathbf{s} \in S_i} \int_{\mathbf{v}} F(d\mathbf{s}|\mathbf{v})G(d\mathbf{v}) \\ &= \sum_{i=1}^N V(F(\cdot|S_i))\mathcal{P}_i^0 = \sum_{i=1}^N \int_{\mathbf{v}} v_i F(d\mathbf{v}|S_i)\mathcal{P}_i^0 = \sum_{i=1}^N \int_{\mathbf{v}} v_i \mathcal{P}_i(\mathbf{v})G(d\mathbf{v}), \end{aligned}$$

where the final step uses the relationship $\Pr(X|Y)\Pr(Y) = \Pr(Y|X)\Pr(X)$.

²⁰For the concavity of entropy, see Cover and Thomas (2006).

The cost of information expressed in (5) is λ times the mutual information between \mathbf{s} and \mathbf{v} . As there is a single posterior leading to each action under the optimal strategy, the joint distribution of the signal and the state is the same as that of the action and the state. Therefore, the cost also equals λ times the mutual information between i and \mathbf{v} . Finally, (9) is implied by the symmetry of mutual information, i.e. for any random variables X and Y , $H(X) - E_Y[H(X|Y)] = H[Y] - E_X[H[Y|X]]$. The cost is thus expressed as the unconditional entropy of i less the expected entropy of i conditional on \mathbf{v} .

Proof of Corollary 1. We begin by showing the solution to Definition 1 solves (10) with a proof by contradiction. Let (F, a) be the optimal strategy, let it attain an expected utility (expected payoff net of information costs) of U_1 , and let \mathcal{P} be the collection induced by it via (6), which does not solve (10). The expected utility from \mathcal{P} in (10) also equals U_1 , see (8) and (9) in Lemma 1.

Since \mathcal{P} does not solve (10), but it satisfies the constraints (11)-(12), there must exist an admissible \mathcal{P}' that provides an expected utility $U_2 > U_1$. Moreover, there must also exist a strategy (F', a') generated from \mathcal{P}' that provides expected utility U_2 , too. For instance, (F', a') can be generated in the following way. Let $\{s_i\}_{i=1}^N$ be N distinct points, i.e. signals, in \mathbb{R}^N , let $F'(s_i, \mathbf{v})$ be the product of $\mathcal{P}_i(\mathbf{v})$ and $G(\mathbf{v})$ and let $a'(F(\mathbf{v}|s_i)) \equiv i$. This strategy clearly satisfies the consistency with G in (4), and it generates utility $U_2 > U_1$ so the strategy (F, a) cannot be optimal.

We now show the solution to (10) solves Definition 1 also by contradiction. Let us assume that \mathcal{P} solves (10) and generates utility U_3 , but it is not induced by any optimal strategy. This means that if F is generated from \mathcal{P} the same way as F' from \mathcal{P}' above, then there exists another admissible strategy \hat{F} that generates a higher utility $U_4 > U_3$. However, this in turn means that $\hat{\mathcal{P}}$ that is induced by \hat{F} also generates expected utility U_4 . This is a contradiction, since \mathcal{P} then cannot be a solution to (10).

PROOF FOR SECTION II

Proof of Theorem 1. The case of $\lambda = 0$ is trivial. When $\lambda > 0$, then the Lagrangian of the problem formulated in Corollary 1 is:

$$\begin{aligned} \mathcal{L}(\mathcal{P}) = & \sum_{i=1}^N \int_{\mathbf{v}} v_i \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}) - \lambda \left(- \sum_{i=1}^N \mathcal{P}_i^0 \log \mathcal{P}_i^0 + \sum_{i=1}^N \int_{\mathbf{v}} \mathcal{P}_i(\mathbf{v}) \log \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}) \right) \\ & + \int_{\mathbf{v}} \xi_i(\mathbf{v}) \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}) - \int_{\mathbf{v}} \mu(\mathbf{v}) \left(\sum_{i=1}^N \mathcal{P}_i(\mathbf{v}) - 1 \right) G(d\mathbf{v}), \end{aligned}$$

where $\xi_i(\mathbf{v}) \geq 0$ are Lagrange multipliers on (11) and $\mu(\mathbf{v})$ are the multipliers on (12).

If $\mathcal{P}_i^0 > 0$, then the first order condition with respect to $\mathcal{P}_i(\mathbf{v})$ is:

$$(B1) \quad v_i + \xi_i(\mathbf{v}) - \mu(\mathbf{v}) + \lambda \left(\log \mathcal{P}_i^0 + 1 - \log \mathcal{P}_i(\mathbf{v}) - 1 \right) = 0.$$

This implies that if $\mathcal{P}_i^0 > 0$ and $v_i > -\infty$, then $\mathcal{P}_i(\mathbf{v}) > 0$ almost surely. To see this, suppose to the contrary that $\mathcal{P}_i(\mathbf{v}) = 0$ on a set of positive measure with respect to G . Since $\xi_i(\mathbf{v}) \geq 0$ and since we also assume that $\mathcal{P}_i^0 > 0$, and thus $\log \mathcal{P}_i^0 > -\infty$, it would have to be $\mu(\mathbf{v})$ going to infinity that would balance $\log \mathcal{P}_i(\mathbf{v}) = -\infty$ to make the first order condition hold. However, if $\mu(\mathbf{v}) = \infty$ on a set of positive measure, then for all such \mathbf{v} in order for (B1) to hold then for all j either $\mathcal{P}_j(\mathbf{v}) = 0$ or $\xi_j(\mathbf{v}) = \infty$. But $\xi_j(\mathbf{v}) > 0$ only if $\mathcal{P}_j(\mathbf{v}) = 0$, when (11) is binding. Therefore, if there exists

i such that $\mathcal{P}_i(\mathbf{v}) = 0$, then $\mathcal{P}_j(\mathbf{v}) = 0$ for all j . This is not possible, since then $\sum_{j=1}^N \mathcal{P}_j(\mathbf{v}) = 0$, which must sum up to 1 and hence a contradiction.

Therefore, whenever \mathcal{P}_i^0 is positive, then the conditional probability $\mathcal{P}_i(\mathbf{v})$ is also positive as long as the realized payoff v_i is not minus infinity. As (11) does not bind, we have $\xi_i(\mathbf{v}) = 0$ and the first order condition can be rearranged to

$$(B2) \quad \mathcal{P}_i(\mathbf{v}) = \mathcal{P}_i^0 e^{(v_i - \mu(\mathbf{v}))/\lambda}.$$

Plugging (B2) into (12), we find:

$$e^{\mu(\mathbf{v})/\lambda} = \sum_i \mathcal{P}_i^0 e^{v_i/\lambda},$$

which we again use in (B2) to arrive at equation (13). Finally, notice that the theorem holds even for $\mathcal{P}_i^0 = 0$, as otherwise equation (7) could not hold.

Proof of Lemma 2. Substitute equation (13) into the objective function to obtain

$$\sum_{i=1}^N \int_{\mathbf{v}} v_i \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}) + \lambda \left\{ \sum_{i=1}^N \mathcal{P}_i^0 \log \mathcal{P}_i^0 - \int_{\mathbf{v}} \left[\sum_{i=1}^N \mathcal{P}_i(\mathbf{v}) \log \left(\frac{\mathcal{P}_i^0 e^{v_i/\lambda}}{\sum_{j=1}^N \mathcal{P}_j^0 e^{v_j/\lambda}} \right) \right] G(d\mathbf{v}) \right\}$$

and rearrange to obtain the new objective function.

Proof of Proposition 1. The solution to the decision maker's problem is unique due to point (2) of Corollary 1, which can be found in online Appendix C.

The decision maker forms a strategy such that $\mathcal{P}_i^0 = 1/N$ for all i . If there were a solution with non-uniform \mathcal{P}_i^0 , then any permutation of the set would necessarily be a solution too, but the solution is unique. Using $\mathcal{P}_i^0 = 1/N$ in equation (13), we arrive at the result.