

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

Filip Matějka* and Alisdair McKay**

November 22, 2011

Abstract

Individuals must often choose among discrete alternatives with imperfect information about their values. Before choosing, they may have an opportunity to study the options, but doing so is costly. This costly information acquisition 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 modified multinomial logit model. The modification arises because the decision maker's prior knowledge and attention allocation strategy affect his evaluation of the alternatives. When the options are *a priori* homogeneous, the standard logit model emerges.

Keywords: discrete choice, information, rational inattention, multinomial logit.

[†] We thank Christian Hellwig, Christopher Sims, Jorgen Weibull, Michael Woodford, Jakub Steiner, Satyajit Chatterjee, Stepan Jurajda, Faruk Gul, Tony Marley, Levent Celik, Andreas Ortmann and Leonidas Spiliopoulos for helpful discussions as well as seminar participants at BU, CERGE, Harvard Institute for Quantitative Social Science, Princeton University and SED 2011.

[‡]This research was funded by GA ČR P402/11/P236 and GDN project RRC 11+001

*CERGE-EI, Prague. filip.matejka@cerge-ei.cz

**Boston University. amckay@bu.edu

1 Introduction

Individuals are frequently confronted with a discrete choice situation in which they do not know the values of the available options, but have an opportunity to investigate the alternatives before making a choice. For example, an employer is able to interview candidates for a job before selecting one to hire. In this context, the decision maker (DM) faces choices of how much to study the options and what to investigate when doing so. For example, the employer might choose how long to spend interviewing the candidates and also choose what questions to ask them during the interview. In most cases, however, it is too costly to investigate the options to the point where their values are known with certainty and, as a result, some uncertainty about the values remains when the choice is made. Because of this uncertainty, the option that is ultimately chosen may not be the one that provides the highest utility to the DM. Moreover, noise in the decision process may lead identical individuals to make different choices. For these reasons, imperfect information naturally leads choices to contain errors and be probabilistic as opposed to deterministic.

In this paper, we explore the behavior of an agent facing a discrete choice when information about the options is costly to acquire and process. In our setting, the DM enters the choice situation with some prior beliefs about the values of the available options. He then processes information about the options in the manner that is optimal given the costs, which we model using the rational inattention framework introduced by Sims (2003, 2006). The major appeal of the rational inattention approach 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 utility-maximizing behavior. After processing information, the DM selects the option that has the highest expected value according to his posterior knowledge.

Our main finding is that the DM chooses probabilistically with choice probabilities that follow a generalized multinomial logit model. If the DM views the options symmetrically *a priori*, then he chooses exactly according to the standard multinomial logit. If the DM's prior knowledge leads him to prefer some options over others, then this prior information is

incorporated into the choice probabilities. In a choice among N options with values v_i for $i \in \{1, \dots, N\}$, the logit model implies that the probability of choosing option i is

$$\frac{e^{v_i/\lambda}}{\sum_{j=1}^N e^{v_j/\lambda}},$$

where λ is a scale parameter. In our model, λ scales the cost of information. Our modified logit formula takes the form

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

where the α_i terms are determined by the DM's prior knowledge of the options and information processing strategy. The DM's information processing strategy is relevant because the DM can focus his attention on those options that he *a priori* believes to be good candidates. The DM's choice of information processing strategy is based on his prior knowledge of the options and, as a result, the α_i terms only reflect the DM's *a priori* beliefs and are not related to the actual values of the options. These adjustments to the logit model can lead the DM to have a systematic positive bias towards an option even when its true value is low. As the cost of information rises, the DM's choice becomes less sensitive to the actual values of the options and more sensitive to his prior beliefs.

The adjustments for the DM's prior knowledge reflect something deeper than just the impact of the prior in standard Bayesian updating. In particular, these adjustments reflect the DM's decisions about how to process information and these decisions affect the choice probabilities in complex ways. We find that this effect breaks the independence of irrelevant alternatives property for which the standard multinomial logit has been criticized.¹ In addition, we find that adding an option to the choice set can increase the probability that an existing option is selected. We also find that changes in the correlation structure of the values of the different options can lead the DM to choose to ignore an option completely even when it may be his best option. These implications are fairly intuitive and are direct consequences of the DM's rational choice of how to allocate his attention.

¹For example, Debreu's (1960) critique, which is now known as the red-bus-blue-bus problem.

The multinomial logit model is perhaps the most commonly used model of discrete choice.² It is so widely used because it is particularly tractable both analytically and computationally and because it has a connection to consumer theory through a random utility model (McFadden, 2001). According to the random utility derivation of the logit, the DM evaluates the options 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 Typically, the randomness in choices is thought to stand for unobserved heterogeneity in tastes across individuals, but sometimes it is attributed to errors of perception.⁴ We provide an explicit foundation for the errors of perception interpretation of the multinomial logit without any distributional assumptions. We find the standard logit model is applicable in some cases, but in other cases the choice probabilities reflect the DM's *a priori* beliefs and attention allocation choices as well as the true values of the options. We show that shifts in the allocation of attention can lead to choice behavior that is inconsistent with any random utility model.

Our work has implications for the interpretation of data on choices. Under the random utility interpretation of the logit model, when one estimates the model, one is estimating the relationship between the typical value, or systematic utility, of an option and covariates that describe different choice situations that arise due to variation in the available options or variation in the characteristics of the individual making the choice. Under our interpretation of the model, an empirical estimate reflects both the values of the available options and the adjustments for prior knowledge and information processing strategies. These adjustments can confound the relationship between values and choice probabilities even when all individuals enter the choice situation with the same prior knowledge of the options.

In addition to empirical applications, the multinomial logit model is often used in applied

²Useful surveys of discrete choice theory and the multinomial logit model are presented by Anderson et al. (1992), McFadden (2001), Train (2009).

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.

⁴See, for example, McFadden (1980, p. S15).

theory.⁵ Future work can build upon our results to study the role of information frictions in these contexts while still exploiting the tractability of the multinomial logit.

The paper is organized as follows: In the remainder of this section we review related work. Section 2 presents the choice setting and discusses the assumptions underlying the rational inattention approach to information frictions. Section 3 characterizes the DM's optimal strategy. Section 4 demonstrates how the DM's prior knowledge influences his choice behavior and establishes that the standard multinomial logit model arises in a situation where the options are symmetric *a priori*. Finally, Section 5 concludes.

Related Literature Our work relates to the literature on rational inattention. Most existing work with rational inattention has focussed on situations where the DM chooses from a continuous choice set.⁶ A few papers, however, consider binary choice problems. Woodford (2009) studies a binary choice of whether to adjust prices, 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 options is continuous. These authors do not make the connection to the multinomial logit model. In an independent paper that is as of yet unfinished, Woodford (2008) notices the connection in the context of a binary choice problem, but does not explore it in further detail.

Our work also relates to studies of discrete choice under imperfect information. Weibull et al. (2007) consider a DM who receives signals about the options before making the choice

⁵The 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 et al. (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.

⁶Rational inattention has mostly been applied in macroeconomic contexts. The major applications have been consumption-savings problems (Sims, 2006; Luo, 2008; Tutino, 2009; Luo and Young, 2009; Maćkowiak and Wiederholt, 2010), price setting (Maćkowiak and Wiederholt, 2009; Matějka, 2010a,b), monetary policy (Paciello and Wiederholt, 2011) and portfolio choices (Van Nieuwerburgh and Veldkamp, 2010; Mondria, 2010).

and allow the DM to select the precision of the signals at a cost. If one assumes that the signals are the true values plus an extreme-value distributed noise term, this becomes a logit model in which the DM can select the scaling factor. Natenzon (2010) has proposed a model in which the DM has Gaussian priors on the utilities of the options and then receives Gaussian signals about the utilities. As information accumulates over time, the DM updates his beliefs about which option is best. A key difference between these models and ours is that Weibull et al. and Natenzon assume particular properties of the noise in agents' observations, which then generate the corresponding properties of the choice probabilities. We derive the properties of the DM's posterior uncertainty endogenously from first principles and show how they can vary depending on the choice set. We find that the generalized logit formula holds for all forms of prior knowledge.

In addition, our work relates to alternative derivations of the multinomial logit. Recently, Gul et al. (2010) have proposed a new axiomatic foundation for the multinomial logit that weakens the Choice Axiom from which Luce (1959) originally derived the logit. While, the Choice Axiom states that the ratio of selection probabilities should remain unchanged as the choice set varies, Gul et al. only require that the ordering of selection probabilities remains unchanged.

The rational inattention approach to information frictions uses information theoretic concepts to measure the amount of information processed by the DM 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). Here we are considering the decision problem of an agent who must acquire information about the values of the alternatives. In this context, the entropy function arises naturally.⁷

⁷In mathematical terms, our work is close to that of Shannon (1959) who derives the multinomial logit formula in an engineering application that is the dual to our problem in the symmetric case. Shannon's question is how quickly 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.

2 The model

In this section, we first describe the agent's decision problem, then we discuss the modeling choices of how the agent processes information.

The DM is presented with a group of N options, from which he must choose one. The values of these options potentially differ and the agent wishes to select the option with the highest value. Let v_i denote the value of option $i \in \{1, \dots, N\}$.

The DM is rationally inattentive in the style of Sims (2003, 2006). He possesses some prior knowledge of the available options; this prior knowledge is described by a joint distribution $G(\mathbf{v})$, where $\mathbf{v} = (v_1, \dots, v_N)$ is the vector of values of the N options. To refine his knowledge, he processes information about the options. This information processing is done through a limited-capacity information channel. One interpretation is that he asks questions about the values and each question comes at a cost.⁸ Finally, the DM chooses the option with the highest expected value.

The DM maximizes the expected value of the selected option minus the cost of information processing. We assume that all information about the N options is available to the DM, but processing the information is costly. If the DM could process information costlessly, he would select the option with the highest value with probability one. With costly information acquisition, the DM must choose the following:

- (i) how much information to process, i.e. how much attention to pay,
- (ii) what pieces of information to process, i.e. what to pay attention to,
- (iii) what option to select conditional on the acquired posterior belief.

Let us first describe how the choice of (i)-(ii) is modeled. In general, what knowledge outcomes can be generated by a specific mechanism of processing information is fully described by a joint distribution of fundamentals and posterior beliefs about them. Blackwell (1953) calls this joint distribution an experiment to emphasize it can be some series of tests

⁸For a textbook treatment of information theory and limited-capacity channels see Cover and Thomas (2006).

that the agent performs in order to gauge the fundamental of interest. The joint distribution describes what the agent can learn from performing the tests.

The novelty of rational inattention is that the DM is allowed to choose the optimal mechanism of processing information considering the cost of acquiring information, (i)-(ii) above. He chooses how to allocate his attention. In Blackwell’s terminology, the DM chooses how to design the experiment. The rationally inattentive DM chooses what pieces of information to process by deciding what questions to ask, what indicators to look at, what media to pay attention to, etc.

The outcome of the DM’s overall strategy, (i)-(iii) above, is a joint distribution of values \mathbf{v} and choices of i . This is because outcomes of (i)-(ii) are described by a joint distribution of \mathbf{v} and posterior beliefs, and because each posterior belief determines a particular choice i from among the options, (iii). It is this joint distribution between \mathbf{v} and i that is the DM’s strategy under the rational inattention approach. It is not necessary to model signals. Nevertheless, the DM’s strategy describes both the choice of how to process information as well as the choice among the options conditional on the posteriors.⁹

Since $G(\mathbf{v})$ is the DM’s prior on the values of all options, which is given, we can describe the DM’s strategy using the conditional probability $\mathcal{P}(i|\mathbf{v}) \in [0, 1]$ for all i and \mathbf{v} . This is the probability of option i being selected when the realized values are \mathbf{v} . Let us denote this probability as $\mathcal{P}_i(\mathbf{v})$.

The DM’s strategy is thus a solution to the following problem:

$$\max_{\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N} \left(\sum_{i=1}^N \int_{\mathbf{v}} v_i \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}) - \text{cost of information processing} \right), \quad (1)$$

⁹Signals do not show up in the final formulation of the problem since each posterior belief is associated with a single i that is selected given that belief. It would not be optimal to select an information structure that would generate two different forms of posterior knowledge leading to the same i , i.e. it would not be optimal to acquire information that is not ultimately used in the choice. This approach is used in Sims (2006) and also in Matějka (2010a), where it is discussed at more length

subject to

$$\sum_{i=1}^N \mathcal{P}_i(\mathbf{v}) = 1 \quad a.s. \quad (2)$$

The first term in (1) is the expected value of the selected option. Finally, we specify the cost of information. Rationally inattentive agents are assumed to process information through channels with a limited information capacity. The cost of information is $\lambda\kappa$, where λ is the unit cost of information and κ the amount of information that the DM processes, which is measured by the expected reduction in the entropy, H , of the distribution representing knowledge of \mathbf{v} . The amount of information processed, κ , is a function of the DM's strategy of how to process information, while λ is a given parameter. This particular form of the cost of information is now common in the literature on rational inattention. In section 2.2, we explain how information theory connects limited-capacity channels to the reduction in entropy.

Entropy is a measure of the uncertainty associated with a random variable. In our case, the random variable is the vector \mathbf{v} and acquiring better knowledge about the values, i.e. narrowing down the belief, is associated with a decrease in the entropy. Sharpening the belief requires processing information. Mathematically, the entropy of a random variable X with a pdf $p(x)$ with respect to a probability measure σ is defined as:

$$H(X) = - \int p(x) \log p(x) \sigma(dx). \quad (3)$$

The expected reduction in the entropy of \mathbf{v} is the difference between the prior entropy of \mathbf{v} and the expectation of the posterior entropy of \mathbf{v} conditional on the chosen option, i . This quantity is also called the mutual information between \mathbf{v} and i . For our purposes, it is convenient to use the symmetry of mutual information and express the amount of information

processed as the expected reduction in the entropy of i conditional on \mathbf{v} :¹⁰

$$\begin{aligned}\kappa(\mathcal{P}, G) &= H(\mathbf{v}) - E_i [H(\mathbf{v}|i)] = H(i) - E_{\mathbf{v}} [H(i|\mathbf{v})] \\ &= - \sum_{i=1}^N \mathcal{P}_i^0 \log \mathcal{P}_i^0 + \int_{\mathbf{v}} \left(\sum_{i=1}^N \mathcal{P}_i(\mathbf{v}) \log \mathcal{P}_i(\mathbf{v}) \right) G(d\mathbf{v}),\end{aligned}\tag{4}$$

where $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ is the collection of conditional probabilities, and \mathcal{P}_i^0 is the unconditional probability of choosing option i ,

$$\mathcal{P}_i^0 = \int_{\mathbf{v}} \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}).$$

We can now state the DM's optimization problem.

Definition 1. *Let $G(\mathbf{v})$ be the DM's prior on the values of a finite number of options and let $\lambda \geq 0$ be the unit cost of information. The discrete choice strategy of the rationally inattentive DM is the collection of conditional probabilities $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ that solves the following optimization problem.*

$$\max_{\mathcal{P}=\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N} \sum_{i=1}^N \int_{\mathbf{v}} v_i \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}) - \lambda \kappa(\mathcal{P}, G),\tag{5}$$

subject to (2), and where $\kappa(\mathcal{P}, G)$ denotes the right hand side of (4).

In the next two subsections, we discuss two of the main modeling assumptions that are at the heart of the literature on rational inattention.

2.1 Endogenous information structure

The DM's optimization problem is formulated as a choice over joint distributions of fundamentals and actions. The only constraint that restricts the DM's choice of distribution is that the conditional probabilities must sum to one for all \mathbf{v} . This means that all the

¹⁰The mutual information between random variables X and Y is $H(X) - E_Y[H(X|Y)]$, which also equals $H(Y) - E_X[H(Y|X)]$, see Cover and Thomas (2006).

information about the N options is available to the DM, so the DM's actions are only truly constrained by the cost of processing the information.

In this formulation, we are not imposing an information structure on the DM. A common approach to modeling information frictions is to posit that there are signals that are drawn from a particular distribution. In some cases, the DM has no choice over the signals. In others, the DM can choose the number or precision of these signals. The rational inattention framework, by contrast, allows the DM to choose all aspects of his posterior uncertainty. In terms of signals, it is as if the DM can choose all features of the distribution that the signals are drawn from.

2.2 Entropy and the cost of information processing

Our results depend crucially on the choice to model the cost of information as λ times the reduction in entropy. Fortunately, entropy provides the exact measure of the cost for rationally inattentive agents who acquire information through a limited-capacity channel. This is a fundamental finding of information theory (Shannon, 1948; Cover and Thomas, 2006).

Using a limited-capacity channel means the DM receives a sequence of symbols (e.g. a list of ones and zeros). The symbols can mean virtually anything: they can represent answers to questions the agent asks, pieces of text or digits he reads, etc. The more information the DM processes, i.e. the more symbols he receives, the tighter his posterior beliefs can be. The capabilities of limited-capacity channels to transmit information are studied in information theory, which is a sub-field of engineering.

Entropy has a foundation in information theory as the exact measure of what information can be passed through channels. The coding theorem of information theory states that any joint distribution of source variables, i.e. fundamentals, and posterior beliefs is achievable by an information channel if and only if the expected decrease in the entropy of the knowledge is less than the amount of information processed, which is proportional to the number of symbols received. Choosing how to process information is then equivalent to choosing how

many questions to ask and what to ask about.¹¹

The assumption that the cost of information processing is $\lambda\kappa$ can be interpreted as saying the cost is proportional to the expected number of questions asked. One could think of the coefficient λ as a shadow cost of allocating attention to this decision problem out of a larger budget of attention that the agent is allocating to many issues. By modeling the cost of information in terms of the number of questions that the DM asks or the number of symbols that he receives, we are modeling a world in which receiving answers to each question with the same number of possible answers is equally costly.¹²

2.3 Existence and uniqueness

Appendix A establishes that a solution to the DM's maximization problem exists. However, the solution may not be unique. For example, if two options always take the same value as one another, then the DM's expected pay-off will not change as probability is shifted between these two options and, as a result, there can be multiple optimal strategies. Appendix A also lays out conditions under which the solution is unique and these conditions require a sufficient independent variation in the values of the options to rule out situations like the one in the example.

¹¹The amount of information per symbol depends on the physical properties of the channel. For instance, if the DM processes information by asking questions with yes or no answers, then the information per question is one bit.

¹²Besides the foundation in information theory, entropy can be derived axiomatically as a natural measure of information. Consider a setting in which there are N possible states of the world each with a certain probability of being the true state. Shannon (1948) asked, how do we measure the amount of information that is communicated when the true state is revealed? He showed that the following axioms imply entropy is the measure of information: i) the measure is increasing in the number of equally likely states, ii) it is continuous in the probabilities of the states, and iii) it is irrelevant if one first learns that the true state lies within some subset and then learns which member of that subset (Shannon, 1948, Theorem 2). The third of these axioms is closely related to the Luce's Choice Axiom (Luce, 1959) which implies the logit model for choice probabilities.

3 Solving the model

In this section, we first derive a general analytical expression for the probability that the DM chooses a particular option conditional on the values of all the options. This expression is not fully explicit as it depends on the unconditional probabilities of choosing each option given by \mathcal{P}_i^0 , for which we still need to solve. We then discuss the solution for the unconditional probabilities.

If information is costless, $\lambda = 0$, then the DM is perfectly attentive and simply selects the option with the highest value with probability one. When $\lambda > 0$, then the Lagrangian of the DM's problem formulated in Definition 1 is:

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

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

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

This can be rearranged to

$$\mathcal{P}_i(\mathbf{v}) = \mathcal{P}_i^0 e^{(v_i - \mu(\mathbf{v})) / \lambda}. \quad (6)$$

Plugging (6) into (2), we find:

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

which we again use in (6) to arrive at the following theorem, which holds even for $\mathcal{P}_i^0 = 0$.

Theorem 1. *If $\lambda > 0$, then the DM forms his strategy such that the conditional choice probabilities satisfy:*

$$\mathcal{P}_i(\mathbf{v}) = \frac{\mathcal{P}_i^0 e^{v_i / \lambda}}{\sum_{j=1}^N \mathcal{P}_j^0 e^{v_j / \lambda}}. \quad (7)$$

If $\lambda = 0$, then the DM selects the option(s) with the highest value with probability one.

From (7) we can understand several properties of the DM's behavior, although it does not fully solve the DM's problem because it includes $\{\mathcal{P}_i^0\}_{i=1}^N$, which are functions of the choice variables.

The unconditional probabilities are by definition independent of a specific realization of the values \mathbf{v} . They are the marginal probabilities of selecting each option 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 , (7) becomes the usual multinomial logit formula. As we discuss in Section 4.1, this happens when G is invariant to permutations of its arguments. In other cases, the conditional choice probabilities are not driven just by $\{e^{v_i/\lambda}\}_{i=1}^N$, as in the logit case, but also by the unconditional probabilities of selecting each option, $\{\mathcal{P}_i^0\}_{i=1}^N$. The effect of the unconditional probabilities is perhaps more obvious if we set $\alpha_i = \lambda \log(\mathcal{P}_i^0)$. α_i reflects the unconditional, i.e. *a priori*, attractiveness of option i . Equation (7) can be rewritten as:

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

Written this way, the selection probabilities can be interpreted as a multinomial logit in which the value of option i is shifted by the term α_i . As the cost of information, λ , rises, the weight on the prior rises, too; the exponents are $v_i/\lambda + \log(\mathcal{P}_i^0)$. The costlier the information is, the less the DM finds out about the realization of \mathbf{v} and the more he decides based on prior knowledge of the options. When an option seems very attractive *a priori*, then it has a relatively high probability of being selected even if its true value is low.

The parameter λ converts bits of information to utils. Therefore, if one scales the values of all of the options by a constant c , while keeping the information cost, λ , fixed, the problem is equivalent to the one with the original values and the information cost scaled by $1/c$. By scaling up the values, one is scaling up the differences between the values and therefore raising the stakes for the DM. The DM chooses to process more information because more is at stake and thus is more likely to select the option that provides the highest utility. The

DM behaves just as he would if the cost of information had fallen.

3.1 Unconditional choice probabilities

Theorem 1 reduces the DM's original problem of selecting $\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$, which are N functions on the support of $G(\mathbf{v})$, to choosing $\{\mathcal{P}_i^0\}_{i=1}^N$, which are N probabilities. We now turn our attention to the solution for these unconditional probabilities. These probabilities must be internally consistent in that $\mathcal{P}_i^0 = \int_{\mathbf{v}} \mathcal{P}_i(\mathbf{v})G(d\mathbf{v})$, where $\mathcal{P}_i(\mathbf{v})$ is given by equation (7). The unconditional probabilities must also maximize the DM's objective function. We can use these conditions to solve for the unconditional probabilities and we demonstrate how to do this in some important cases in the next section. Analytical solutions are possible in the problems we study in this paper. In most other cases, the best approach is to numerically maximize the DM's objective with respect to the unconditional choice probabilities. Before turning to these important cases, we establish some basic properties of the solution.

- Proposition 1.**
1. *If $\lambda = 0$, then the DM always chooses the option(s) with the highest value. The prior does not influence conditional choice probabilities.*
 2. *If $\lambda = \infty$, or any time the DM decides not to process information, then the DM always chooses the option(s) with the highest expected value with respect to the prior knowledge G only. Choice probabilities do not depend on realized values.*
 3. *If an option j is dominated by another option k , i.e. $v_j < v_k$ with probability one, then the DM never chooses option j , i.e. $\mathcal{P}_j^0 = 0$.*
 4. *If the value of option k is increased by $\omega > 0$ in all states of the world, the unconditional probability \mathcal{P}_k^0 does not decrease.*

Proof: Appendix B.1.

The unconditional probabilities depend on the whole prior distribution $G(\mathbf{v})$. Bayesian updating has the effect that higher values of an option in some states of the world increase the probability that the option is selected in other states. Moreover, changes in the prior

distribution can lead the DM to allocate his attention differently, making him more or less likely to select a given option. In the next section we study how the joint distribution of the DM's *a priori* knowledge of the options in the choice set affects the choice probabilities.

4 The influence of prior knowledge of the choice set

The difference between the behavior of the rationally inattentive agent and the standard logit model comes from the presence of the unconditional choice probabilities in equation (7). In this section, we show how these probabilities are determined in series of different contexts.

We begin in Section 4.1 with a general result: when the prior is symmetric so that the options are exchangeable in the prior, the model reduces to the standard logit model. When the prior is not symmetric, however, the DM may have some prior knowledge of the options that will influence his choice. Section 4.2 presents the intuition for how the solution behaves when the options are not symmetric *a priori*. We then consider a context in which the logit model has been criticized, namely when two options are duplicates and show the rationally inattentive DM treats duplicate options as a single option. Finally, we demonstrate that the model can generate behavior that is inconsistent with random utility maximization.

4.1 A priori homogeneous options: the multinomial logit

Let us assume that all the options seem identical to the DM *a priori* and are exchangeable in the prior G . We call the options *a priori homogeneous* if and only if $G(\mathbf{v})$ is invariant to all permutations of the entries of \mathbf{v} .

Problem 1. *The DM chooses $i \in \{1, \dots, N\}$, where the options are a priori homogeneous and take different values with positive probability.*

Theorem 2. *In Problem 1, the probability of choosing option i as a function of the realized values of all of options is:*

$$\mathcal{P}_i(\mathbf{v}) = \frac{e^{v_i/\lambda}}{\sum_{j=1}^N e^{v_j/\lambda}}, \quad (8)$$

which is exactly the multinomial logit formula.

Proof: Appendix B.2. We show that the homogeneity of options implies that the unconditional probabilities are uniform, (7) then takes the form of the logit. The assumption on the difference of values is needed for uniqueness. If all options are always the same, then whatever choice the DM makes, the realized value is independent of it, thus the solution is not unique.

Let us emphasize that $\mathcal{P}_i(\mathbf{v})$ does not depend on the prior G . As long as the options are *a priori* homogeneous, the resulting choice probabilities take the form of (8). This feature is particularly useful as it makes applications of the rational inattention framework very simple in this case.

The DM always chooses to process some information, which is not necessarily the case when the prior is asymmetric. Here the marginal expected value of additional information is initially infinite and then decreases as the DM processes more information. Therefore, the DM chooses to process some positive amount of information as long as λ is finite. The marginal value of information is initially high because the uniform distribution of choice probabilities maximizes the entropy and, as a result, the derivative of the entropy reduction with respect to the choice probabilities is zero at this point.

4.2 Departure from logit

We now explore an example illustrating how the DM's prior influences the choice probabilities. There are two options, one of which has a known value while the other takes one of two values. One interpretation is that the known option is an outside option or reservation value. This problem is a simple benchmark that exhibits the basic features of most solutions to problems with asymmetric priors.

Problem 2. *The DM chooses $i \in \{1, 2\}$. The value of option 1 is distributed as $v_1 = 0$ with the probability g_0 and $v_1 = 1$ with the probability $1 - g_0$. Option 2 carries the value $V_2 = R \in (0, 1)$ with certainty.*




Figure 1: \mathcal{P}_1^0 as a function of R and $\lambda = 0.1, g_0 = 0.5$.

To solve the problem, we must find $\{\mathcal{P}_i^0\}_{i=1}^2$. We show in Appendix C.1 that the solution is:

$$\mathcal{P}_1^0 = \max \left(0, \min \left(1, -\frac{e^{\frac{R}{\lambda}} \left(-e^{\frac{1}{\lambda}} + e^{\frac{R}{\lambda}} - g_0 + g_0 e^{\frac{1}{\lambda}} \right)}{\left(e^{\frac{1}{\lambda}} - e^{\frac{R}{\lambda}} \right) \left(-1 + e^{\frac{R}{\lambda}} \right)} \right) \right) \right) \quad (9)$$

$$\mathcal{P}_2^0 = 1 - \mathcal{P}_1^0.$$

For a given set of parameters, the unconditional probability \mathcal{P}_1^0 as a function of R is shown in Figure 1. For R close to 0 or to 1, the DM decides not to process information and selects one of the options with certainty. In the middle range however, the DM does process information and the selection of option 1 is less and less probable as the reservation value, R , increases, since option 2 is more and more appealing. For $g_0 = 1/2$ and $R = 1/2$, solutions take the form of the multinomial logit, i.e. $\mathcal{P}_1^0 = \mathcal{P}_2^0 = 1/2$. If the DM observed the values, he would choose option 1 with the probability $(1 - g_0) = 1/2$ for any reservation value R . However, the rationally inattentive agent chooses option 1 with higher probability when R is low.

Figure 2 again shows the dependance on R , but this time it presents the probability of selecting the first option *conditional* on the realized value $v_1 = 1$, it is $\mathcal{P}_1(1, R)$. Since $R < 1$, it would be optimal to always select the option 1 when its value is 1. The DM obviously does not choose to do that because he is not sure what the realized value is. When R is high, the DM processes less information and selects a low \mathcal{P}_1^0 . As a result, $\mathcal{P}_1(1, R)$ is low.




Figure 2: $\mathcal{P}_1^0(1, R)$ as a function of R and $\lambda = 0.1, g_0 = 0.5$.




Figure 3: \mathcal{P}_1^0 as a function of λ evaluated at various values of g_0 and $R = 0.5$.

In general, one would expect that as R increases, the DM would be more willing to reject option 1 and receive the certain value R . Indeed, differentiating the non-constant part of (9) one finds that the function is non-increasing in R . Similarly, the unconditional probability of selecting option 1 falls as g_0 rises, as it is more likely to have a low value. Moreover, we see from equation (9) that, for $R \in (0, 1)$, \mathcal{P}_1^0 equals 1 for g_0 in some neighborhood of 0 and it equals 0 for g_0 close to 1.¹³ For these parameters, the DM chooses not to process information.

The following Proposition summarizes the immediate implications of equation (9). Moreover, the findings hold for any values of the uncertain option $\{a, b\}$ such that $R \in (a, b)$.

Proposition 2. *Solutions to Problem 2 have the following properties:*

1. *The unconditional probability of option 1, \mathcal{P}_1^0 , is a non-increasing function of g_0 and the value R of the other option.*
2. *For all $R \in (0, 1)$ and $\lambda > 0$, there exist g_m and g_M in $(0, 1)$ such that if $g_0 \leq g_m$, the DM does not process any information and selects option 1 with probability one. Similarly, if $g_0 \geq g_M$, the DM processes no information and selects option 2 with probability one.*

Figure 3 plots \mathcal{P}_1^0 as a function of the information cost λ for three values of the prior, g_0 . When $\lambda = 0$, \mathcal{P}_1^0 is just equal to $1 - g_0$ because the DM will have perfect knowledge of the value of option 1 and choose it when it has a high value, which occurs with probability $1 - g_0$. As λ increases, \mathcal{P}_1^0 fans out away from $1 - g_0$ because the DM no longer possesses perfect knowledge about the value of option 1 and eventually just selects the option with the higher expected value according to the prior.

4.3 Duplicates and independence from irrelevant alternatives

The multinomial logit has well known difficulties when some options are similar or duplicates. The difficulties stem from the property of independence from irrelevant alternatives (IIA),

¹³The non-constant argument on the right-hand side of (9) is continuous and decreasing in g_0 , and it is greater than 1 at $g_0 = 0$ and negative at $g_0 = 1$.

which states that the ratio of the choice probabilities for two alternatives is independent of what other alternatives are included in the choice set. Debreu (1960) criticized this property for being counter-intuitive. The well known example goes: The DM 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 DM is indifferent to the color of the bus, then IIA—and therefore the multinomial logit, which can be derived from IIA—implies probabilities of $1/3, 1/3, 1/3$. Debreu argued that this is counter-intuitive because duplicating one option should not materially change the choice problem.

The behavior of the rationally inattentive agent does not satisfy IIA and as a result is not subject to Debreu’s critique. IIA does not need to hold since the unconditional choice probabilities can change in complex ways as new choices are added to the set of available alternatives and these changes push the choice probabilities away from the logit.

We formalize the notion of duplicate options as a scenario in which two options have perfectly correlated values across different states of the world. In this section, we begin by showing that the rationally inattentive DM treats duplicate options as a single option. We then extend this idea to consider other implications of the correlation structure of the values of the options.

4.3.1 Duplicates

We study a generalized version of Debreu’s bus problem to analyze how the rationally inattentive agent treats duplicate options. In our framework, we define duplicates as options that carry the same value in all states of the world although this common value may be unknown.

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

Problem 3. *The DM chooses $i \in \{1, \dots, N + 1\}$, where the options N and $N + 1$ are duplicates.*

The following theorem states that duplicate options are treated as a single option. We compare the choice probabilities in two choice problems, where the second one is constructed from the first by duplicating one option. In the first problem, the DM's prior is $G(\mathbf{v})$, where $\mathbf{v} \in \mathbb{R}^N$. In the second problem, the DM's prior is $\hat{G}(\mathbf{u})$, where $\mathbf{u} \in \mathbb{R}^{N+1}$. \hat{G} is generated from G by duplicating option N . This means that options 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} .

Theorem 3. *If $\{\mathcal{P}_i^0\}_{i=1}^N$ and $\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ are unconditional and conditional choice probabilities corresponding to a solution to Problem 3, then $\{\hat{\mathcal{P}}_i(\mathbf{u})\}_{i=1}^{N+1}$ solve the corresponding problem with the added duplicate of the option N if and only they satisfy the following:*

$$\hat{\mathcal{P}}_i(\mathbf{u}) = \mathcal{P}_i(\mathbf{v}), \quad \forall i < N \tag{10}$$

$$\hat{\mathcal{P}}_N(\mathbf{u}) + \hat{\mathcal{P}}_{N+1}(\mathbf{u}) = \mathcal{P}_N(\mathbf{v}), \tag{11}$$

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: Appendix B.3. The implication of this theorem is that the DM treats duplicate options as though they were a single option.

4.3.2 Correlated values and attention allocation

The rationally inattentive agent does not just collect exogenously given signals, but decides how to process information. If the options are not homogeneous, the DM can choose to investigate different options in different levels of detail. We now explore a choice among three options, where two options have positively or negatively correlated values. Even though all three options have the same *a priori* expected value, in some cases the DM will ignore one of the options completely.

Problem 4. *The DM chooses from the set {red bus, blue bus, train}. The DM knows the value of the train exactly, $v_t = 1/2$. The buses each take one of two values, either 0 or 1,*




Figure 4: 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.

with expected values $1/2$ for each, the correlation coefficient between their values is ρ . The joint distribution of the values of all three options is:

$$\begin{aligned}
 g(0, 0, 1/2) &= \frac{1}{4}(1 + \rho) \\
 g(1, 0, 1/2) &= \frac{1}{4}(1 - \rho) \\
 g(0, 1, 1/2) &= \frac{1}{4}(1 - \rho) \\
 g(1, 1, 1/2) &= \frac{1}{4}(1 + \rho).
 \end{aligned} \tag{12}$$

In Appendix C.2 we describe how to solve the problem analytically. Figure 4 illustrates the behavior of the model for various values of ρ and λ . The figure shows the unconditional probability that the DM selects a bus of a given color (the probability is the same for both buses). As the correlation between the values of the buses decreases, the probability that a bus carries the largest value among the three options increases and the unconditional probability of choosing either bus increases, too. If the buses' values are perfectly correlated, then the sum of their probabilities is 0.5, they are effectively treated as one option, 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 DM economizes on information by paying more attention to choosing among the buses and less to assessing their values relative to the reservation value $1/2$.

The choice probabilities strongly reflect the endogeneity of the information structure in this case. As the correlation decreases, the DM knows that the best option is more likely to be one of the buses. As a result, the DM 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 is actually the best option.

4.4 Relation to random utility models

The standard multinomial logit model, with its IIA property, has the feature that adding another option to the choice set reduces the choice probabilities of existing options in a proportionate manner.¹⁴ The same is not true of our generalized logit model because the unconditional choice probabilities depend on the full choice set. In fact, adding an additional option can even raise the probability that an existing option is selected. This type of behavior is not just inconsistent with the standard logit model, but is inconsistent with any random utility model. We now demonstrate this possibility with an example.

Problem 5. *Suppose there are three options and two states of the world. The options take the following values in the two states of the world*

	<i>state 1</i>	<i>state 2</i>
<i>option 1</i>	0	1
<i>option 2</i>	1/2	1/2
<i>option 3</i>	Y	$-Y$

States 1 and 2 have prior probabilities $g(1)$ and $g(2)$, respectively.

First, consider a variant of this choice situation in which only options 1 and 2 are available.

¹⁴Adding option $N + 1$ to the choice set reduces the probability of option $i \in \{1, \dots, N\}$ by a factor of $\frac{\sum_{j=1}^N \exp(v_j/\lambda)}{\sum_{j=1}^{N+1} \exp(v_j/\lambda)}$.

Using our results from Problem 2, we know that there exists $g(1) \in (0, 1)$ large enough that the DM will not process information and select option 2 with probability 1 in all states of the world so $\mathcal{P}_1^0 = 0$. Now add option 3 to the choice set. For a large enough value of Y and $g(1) \in (0, 1)$, the DM will find it worthwhile to process information about the state of the world in order to determine whether option 3 should be selected. Given that the DM will now have information about the state of the world, if state 2 is realized, the DM might as well select option 1. From an *a priori* perspective, there is a positive probability of selecting option 1 so $\mathcal{P}_1^0 > 0$. The choice probabilities conditional on the realization of the state of the world are given by equation (7), which implies that the probability of selecting option 1 is zero if $\mathcal{P}_1^0 = 0$ and positive if $\mathcal{P}_1^0 > 0$ and all options have finite values. So we have the following.

Proposition 3. *For $\lambda > 0$, there exist $g(1) \in (0, 1)$ and $Y > 0$ such that adding option 3 to the choice set in Problem 5 increases the probability that option 1 is selected in all states of the world.*

Proof: Appendix C.3.

Corollary 1. *The behavior of a rationally inattentive agent cannot always be described by a random utility model.*

Proof. Random utility models obey a regularity condition: the probability of selecting a given option cannot be increased by expanding the choice set (Luce and Suppes, 1965, p. 342). □

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.

5 Conclusion

In this paper, we have studied the optimal behavior of a rationally inattentive agent who faces a discrete choice problem. This model gives rise to a version of the multinomial logit

model. This result is derived from assumptions about the technology that the agent uses to process information and is not driven by specific assumptions concerning the kinds of signals the agent acquires.

The behavior of the rationally inattentive agent differs from the standard logit model in that the values of the available options are adjusted to reflect the DM's *a priori* beliefs and information processing decisions. When the agent views the options as symmetric or interchangeable *a priori* this adjustment is the same for all of the options and the model reduces to the standard logit model.

An implication of the relationship between rational inattention and the multinomial logit model is that future work can incorporate rational inattention into larger models that involve discrete choices subject to information frictions by exploiting the tractability of the multinomial logit.

References

- Anas, A. (1983). Discrete choice theory, information theory and the multinomial logit and gravity models. *Transportation Research Part B: Methodological*, 17(1):13–23.
- Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.
- Blackwell, D. (1953). Equivalent comparison of experiments. *Annals of Mathematical Statistics*, (24).
- Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. Wiley, Hoboken, NJ.
- Debreu, G. (1960). Review of individual choice behavior by R. D. Luce. *American Economic Review*, 50(1).
- Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the u.s. automobile industry. *Econometrica*, 63(4).
- Gul, F., Natenzon, P., and Pesendorfer, W. (2010). Random choice as behavioral optimization. Princeton University Working Paper.
- Lindbeck, A. and Weibull, J. W. (1987). Balanced-budget redistribution as the outcome of political competition. *Public Choice*, 52.
- Luce, R. D. (1959). *Individual Choice Behavior: a Theoretical Analysis*. Wiley, New York.
- Luce, R. D. and Suppes, P. (1965). Preference, utility, and subjective probability. In Luce, R. D.; Bush, R. and Galanter, E., editors, *Handbook of Mathematical Psychology*, volume 3, pages 249–410. Wiley, New York.
- Luo, Y. (2008). Consumption dynamics under information processing constraints. *Review of Economic Dynamics*, 11(2).

- Luo, Y. and Young, E. R. (2009). Rational inattention and aggregate fluctuations. *The B.E. Journal of Macroeconomics*, 9(1).
- Mackowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *The American Economic Review*, 99.
- Maćkowiak, B. A. and Wiederholt, M. (2010). Business cycle dynamics under rational inattention. CEPR Discussion Papers 7691, C.E.P.R. Discussion Papers.
- Mattsson, L.-G. and Weibull, J. W. (2002). Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior*, 41(1).
- Matějka, F. (2010a). Rationally inattentive seller: Sales and discrete pricing. CERGE-EI Working Papers wp408.
- Matějka, F. (2010b). Rigid pricing and rationally inattentive consumer. CERGE-EI Working Papers wp409.
- Matějka, F. and Sims, C. A. (2010). Discrete actions in information-constrained tracking problems. Technical report, Princeton University.
- McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*. Academic Press, New York.
- McFadden, D. (1980). Econometric models for probabilistic choice among products. *The Journal of Business*, 53(3).
- McFadden, D. (2001). Economic choices. *The American Economic Review*, 91(3):pp. 351–378.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1).
- Mondria, J. (2010). Portfolio choice, attention allocation, and price comovement. *Journal of Economic Theory*, 145(5).

- Natenzon, P. (2010). Random choice and learning. Working paper, Princeton University.
- Paciello, L. and Wiederholt, M. (2011). Exogenous information, endogenous information and optimal monetary policy. Working paper, Northwestern University.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27.
- Shannon, C. E. (1959). Coding theorems for a discrete source with a fidelity criterion. In *Institute of Radio Engineers, International Convention Record*, volume 7.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3).
- Sims, C. 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).
- Train, K. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press, Cambridge, U.K.
- Tutino, A. (2009). The rigidity of choice: Lifetime savings under information-processing constraints. Working paper.
- Van Nieuwerburgh, S. and Veldkamp, L. (2010). Information acquisition and under-diversification. *Review of Economic Studies*, 77(2).
- Verboven, F. (1996). International price discrimination in the european car market. *RAND Journal of Economics*, 27(2).
- Weibull, J. W., Mattsson, L.-G., and Voorneveld, M. (2007). Better may be worse: Some monotonicity results and paradoxes in discrete choice under uncertainty. *Theory and Decision*, 63.

- Woodford, M. (2008). Inattention as a source of randomized discrete adjustment.
- Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics*, 56(Supplement 1).
- Yang, M. (2011). Coordination with rational inattention. Princeton University Working Paper.
- Yellott, J. 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).

A Existence and Uniqueness

Lemma 1. *The DM's optimization problem in Definition 1 always has a solution.*

Proof: Since (7) is a necessary condition for the maximum, then the collection $\{\mathcal{P}_i^0\}_{i=1}^N$ determines the whole solution. However, the objective is a continuous function of $\{\mathcal{P}_i^0\}_{i=1}^N$, since $\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ is also a continuous function of $\{\mathcal{P}_i^0\}_{i=1}^N$. Moreover, the admissible set for $\{\mathcal{P}_i^0\}_{i=1}^N$ is compact. Therefore, the maximum always exists. \square

Assessing uniqueness is less straightforward. The following assumptions are each sufficient conditions for a unique solution to exist, which we prove below.

Assumption 1. *The prior $G(\mathbf{v})$ is invariant with respect to permutations of the entries of \mathbf{v} . Moreover, v_i and v_j are not almost surely equal for all $i \neq j$.*

Assumption 2. *$N = 2$, and the values of the two options are not almost surely equal.*

Assumption 3. *For all but at most one $k \in \{1, \dots, N\}$, there exist two sets $S_1 \subset \mathbb{R}^N, S_2 \subset \mathbb{R}^N$ with positive probability measures with respect to the prior, $G(\mathbf{v})$, such that for all $\mathbf{v}_1 \in S_1$ there exists $\mathbf{v}_2 \in S_2$ where \mathbf{v}_1 and \mathbf{v}_2 differ in k^{th} entry only.*

In Assumption 1, the condition that the options are exchangeable in the prior is a formalization of the notion that they are viewed symmetrically *ex ante*. In Section 4.1 we call such options *a priori* homogeneous. The second part of the assumption is that there is some positive probability that the options have different values. When $N = 2$, as in Assumption 2, the solution is unique even if the prior is not symmetric. For $N > 2$ and *ex ante* asymmetric options, we have Assumption 3. In words, this assumption says that there is independent variation in the value of all options except possibly one. Assumption 3 is, for instance, satisfied if the values of the options are independently distributed and no more than one of their marginals is degenerate to a single point. Although the assumption is quite a bit weaker than independence as it just requires that there is not some form of perfect co-movement between the values.

We start by proving a lemma that we then use in proving the following uniqueness results.

Lemma 2. *If $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ and $\hat{\mathcal{P}} = \{\hat{\mathcal{P}}_i(\mathbf{v})\}_{i=1}^N$ are two distinct solutions to the DM's optimization problem, then the unconditional probabilities satisfy*

$$\sum_{i=1}^N (\mathcal{P}_i^0 - \hat{\mathcal{P}}_i^0) e^{v_i/\lambda} = 0 \quad a.s. \quad (13)$$

Proof: Mutual information is a convex function of the joint distribution of the two variables.¹⁵ The objective (5) is thus a concave functional: the first term is linear and the second is concave. Moreover, the admissible set of $\{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$, satisfying the constraints is convex. Therefore, any convex linear combination $\tilde{\mathcal{P}}(\xi)$ of the solutions \mathcal{P} and $\hat{\mathcal{P}}$ is also a solution. Notice that the unconditional probabilities of the convex combination will satisfy:

$$\tilde{\mathcal{P}}_i^0(\xi) = \mathcal{P}_i^0 + \xi (\hat{\mathcal{P}}_i^0 - \mathcal{P}_i^0) \quad \xi \in [0, 1], \forall i. \quad (14)$$

Since the convex combination is also a solution, it must satisfy (7), which is necessary. Therefore, we have

$$\tilde{\mathcal{P}}_i(\mathbf{v}; \xi) = \frac{\tilde{\mathcal{P}}_i^0(\xi) e^{v_i/\lambda}}{\sum_{j=1}^N \tilde{\mathcal{P}}_j^0(\xi) e^{v_j/\lambda}}.$$

$\tilde{\mathcal{P}}_i(\mathbf{v}; \xi)/\tilde{\mathcal{P}}_i^0(\xi)$ gives the conditional distribution over \mathbf{v} conditional on option i being selected. This must integrate to one:

$$\int \frac{e^{v_i/\lambda}}{\sum_{j=1}^N \tilde{\mathcal{P}}_j^0(\xi) e^{v_j/\lambda}} G(d\mathbf{v}) = 1 \quad \text{for all } \xi \in [0, 1]. \quad (15)$$

Let us express the second derivative of the left hand side of (15) with respect to ξ at $\xi = 0$. The derivative has to equal zero in order for (15) to hold for all $\xi \in [0, 1]$. It equals:

$$\int \frac{e^{v_i/\lambda} \left(\sum_{j=1}^N (\hat{\mathcal{P}}_j^0 - \mathcal{P}_j^0) e^{v_j/\lambda} \right)^2}{\left(\sum_{j=1}^N \tilde{\mathcal{P}}_j^0(\xi) e^{v_j/\lambda} \right)^3} G(d\mathbf{v}). \quad (16)$$

¹⁵See Chapter 2 in Cover and Thomas (2006).

Therefore, for the two different solutions to exist, (13) has to hold. \square

Lemma 3. *If any of Assumptions 1, 2, or 3 holds, then the solution to the DM's optimization problem is unique.*

Proof: We now show that each of Assumptions 1-3 is sufficient for uniqueness. We start with *Assumption 1*: Let us assume the solution is not unique. There exist two different solutions, $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ and $\hat{\mathcal{P}} = \{\hat{\mathcal{P}}_i(\mathbf{v})\}_{i=1}^N$. As (7) is necessary for a solution, it follows that if the solutions have the same unconditional probabilities, $\{\mathcal{P}_i^0\}_{i=1}^N$ and $\{\hat{\mathcal{P}}_i^0\}_{i=1}^N$, then they must be the same solution. Therefore, the solutions must have different unconditional probabilities in order to be distinct. In addition, Lemma 2 establishes that (13) must hold.

According to Assumption 1, the options have different values with positive probability and the prior is invariant to all permutations, there exists $S_1 \subset \mathbb{R}^N$, a set of value vectors \mathbf{v} , of a positive probability measure w.r.t. G such that $v_1 \neq v_2$ for all $\mathbf{v} \in S_1$. Let $S_2 \subset \mathbb{R}^N$ be another set of vectors that is generated from S_1 by switching v_1 and v_2 of all its vectors. Therefore, S_2 also has a positive measure, in fact it has the same measure as S_1 .

The solution is not unique, therefore (13) holds almost surely in both S_1 and S_2 . Since S_2 was generated from S_1 by switching the first two entries of \mathbf{v} , (13) can be expressed in terms of almost sure equality on S_1 for both sets:

$$(\mathcal{P}_1^0 - \hat{\mathcal{P}}_1^0)e^{v_1/\lambda} + (\mathcal{P}_2^0 - \hat{\mathcal{P}}_2^0)e^{v_2/\lambda} + \sum_{i=3}^N (\mathcal{P}_i^0 - \hat{\mathcal{P}}_i^0)e^{v_i/\lambda} = 0 \quad a.s. \text{ in } S_1 \quad (17)$$

$$(\mathcal{P}_1^0 - \hat{\mathcal{P}}_1^0)e^{v_2/\lambda} + (\mathcal{P}_2^0 - \hat{\mathcal{P}}_2^0)e^{v_1/\lambda} + \sum_{i=3}^N (\mathcal{P}_i^0 - \hat{\mathcal{P}}_i^0)e^{v_i/\lambda} = 0 \quad a.s. \text{ in } S_1 \quad (18)$$

Subtracting the two, we get:

$$(\Delta_1 - \Delta_2) \left(e^{\frac{v_1}{\lambda}} - e^{\frac{v_2}{\lambda}} \right) = 0 \quad a.s. \text{ in } S_1, \quad (19)$$

where Δ_i denotes $(\mathcal{P}_i^0 - \hat{\mathcal{P}}_i^0)$. This implies that $\Delta_1 = \Delta_2$, because S_1 is of positive measure and $v_1 \neq v_2$. However, since the prior is invariant to permutations, we could repeat this argument for other pairs of options so all Δ_i for i in $\{1, \dots, N\}$ are equal to each other. Let

us denote the quantity they are equal to as Δ . Because $1 = \sum_{i=1}^N \mathcal{P}_i^0 = \sum_{i=1}^N \hat{\mathcal{P}}_i^0$, then

$$\sum_{i=1}^N (\mathcal{P}_i^0 - \hat{\mathcal{P}}_i^0) = \sum_{i=1}^N \Delta_i = N\Delta = 0.$$

This implies that $\Delta = 0$, unconditional probabilities of the two solutions are equal. Therefore, we find that the solution must be unique.

Assumption 2: for $N = 2$, $(\mathcal{P}_1^0 - \hat{\mathcal{P}}_1^0) = -(\mathcal{P}_2^0 - \hat{\mathcal{P}}_2^0)$, since the unconditional probabilities of each solution sum up to one. (13) takes the form:

$$(\mathcal{P}_1^0 - \hat{\mathcal{P}}_1^0)(e^{\frac{v_1}{\lambda}} - e^{\frac{v_2}{\lambda}}) = 0 \quad a.s. \quad (20)$$

If v_1 and v_2 are not equal almost surely, then $\mathcal{P}_1^0 = \hat{\mathcal{P}}_1^0$. The solution must be unique.

Assumption 3: If there are two different solutions with $\{\mathcal{P}_i^0\}_{i=1}^N$ and $\{\hat{\mathcal{P}}_i^0\}_{i=1}^N$, then there exist k_1, k_2 such that $(\mathcal{P}_k^0 - \hat{\mathcal{P}}_k^0) \neq 0$ for $k \in \{k_1, k_2\}$, since probabilities sum up to one. In other words, if the two solutions differ, then they have to differ at least in two entries. Using Assumption 3, we know that at least one of k_1, k_2 above satisfies the second part of Assumption 3. Let it be k_1 . (13) then implies

$$(\mathcal{P}_{k_1}^0 - \hat{\mathcal{P}}_{k_1}^0)e^{\frac{v_{k_1}}{\lambda}} = 0 \quad a.s. \quad (21)$$

The solution must be unique. □

Perhaps an illustrative interpretation of non-unique solutions is that there always exist options that can be eliminated for the DM to still achieve the same expected utility. This process can be repeated until the solution is unique. In other words, there exist options that the DM can ignore and never select. If two solutions exist, then all their linear combinations are solutions too as long as all \mathcal{P}_i^0 's are non-negative. Therefore, there exists k , such that a solution with $\mathcal{P}_k^0 = 0$ exists. The simplest example is the duplicates, see Problem 3 in Section 4.3, where either of the two options can be eliminated.

B Main proofs

B.1 Proposition 1

Statements 1 and 2 are trivial. With $\lambda = 0$, the information constraint is not binding, while with $\lambda = \infty$ the DM does not process any information and must make his decision based on the prior belief alone.

Proof of statement 3: by contradiction. Option j is dominated by option k . Let us assume that $\mathcal{P}_j^0 > 0$, where \mathcal{P} is the DM's optimal strategy. We show there exists another strategy $\hat{\mathcal{P}}$ such that $\hat{\mathcal{P}}_j^0 = 0$ that generates higher expected utility than \mathcal{P} .

Let $\hat{\mathcal{P}}$ be generated from \mathcal{P} in the following way:

$$\hat{\mathcal{P}}_i(\mathbf{v}) = \mathcal{P}_i(\mathbf{v}) \quad \forall i \neq j, k \quad (22)$$

$$\hat{\mathcal{P}}_j(\mathbf{v}) = 0 \quad (23)$$

$$\hat{\mathcal{P}}_k(\mathbf{v}) = \mathcal{P}_j(\mathbf{v}) + \mathcal{P}_k(\mathbf{v}). \quad (24)$$

$\hat{\mathcal{P}}$ is constructed from \mathcal{P} by relocating the probability distribution conditional on j onto the distribution conditional on k . This strategy certainly generates a higher expected value of the selected option, since k dominates j .

Moreover, the cost of information of $\hat{\mathcal{P}}$ is not higher than that of \mathcal{P} . The amount of information processed is the difference between the entropy of the prior and the expected entropy of posteriors, $H(\mathbf{v}) - E_i[H(\mathbf{v}|i)]$, see equation (4).¹⁶ The prior entropy is fixed, given by G . Since entropy is a concave function of the distribution, and the posterior conditional on k of $\hat{\mathcal{P}}$ is the sum of the posteriors of \mathcal{P} conditional on j and k , then the expected entropy of the posteriors of $\hat{\mathcal{P}}$ is not lower than that of \mathcal{P} . The original strategy \mathcal{P} requires at least as much information as $\hat{\mathcal{P}}$, and thus $\hat{\mathcal{P}}$ generates higher expected utility. \mathcal{P}_j^0 must equal zero. □

¹⁶In this proof, it is convenient to use the symmetry of mutual information, equation (4), and express it in terms of entropies of the prior and posteriors, rather than using the RHS of (4). Therefore, while in most parts of the text we use the probability of i conditional on \mathbf{v} , which is $\mathcal{P}_i(\mathbf{v})$, here the expectation of the entropy of the posterior is a function of the conditional distributions of \mathbf{v} conditional on different values of i .

Proof of statement 4: by contradiction. Let us assume that

$$\mathcal{P}_k^0 > \hat{\mathcal{P}}_k^0, \quad (25)$$

where $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ is a solution to an original problem with the prior G and $\hat{\mathcal{P}} = \{\hat{\mathcal{P}}_i(\mathbf{v})\}_{i=1}^N$ is a solution to the problem with \hat{G} , which is generated from G such that the value of option k is increased by $\omega > 0$:

$$\hat{G}(v_1, \dots, v_k, \dots, v_N) := G(v_1, \dots, v_k - \omega, \dots, v_N). \quad (26)$$

Let $U(\mathcal{P}, G)$ stand for the DM's expected utility derived from the strategy $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ with the prior G ,

$$U(\mathcal{P}, G) = \sum_{i=1}^N \int_{\mathbf{v}} v_i \mathcal{P}_i(\mathbf{v}) G(d\mathbf{v}) - \lambda \kappa(\mathcal{P}, G), \quad (27)$$

where $\kappa(\mathcal{P}, G)$ is given by (4). The following inequalities are statements that \mathcal{P} and $\hat{\mathcal{P}}$ are solutions to the two problems.

$$U(\mathcal{P}, G) \geq U(\mathcal{P}^*, G) \quad (28)$$

$$U(\hat{\mathcal{P}}, \hat{G}) \geq U(\hat{\mathcal{P}}^*, \hat{G}), \quad (29)$$

where \mathcal{P}^* is any collection of conditional distributions. Let the collection $\mathcal{P}' = \{\mathcal{P}'_i(\mathbf{v})\}_{i=1}^N$ be generated from $\mathcal{P} = \{\mathcal{P}_i(\mathbf{v})\}_{i=1}^N$ and let $\hat{\mathcal{P}}'$ be generated from $\hat{\mathcal{P}}$ such that

$$\mathcal{P}'_i(v_1, \dots, v_k, \dots, v_N) = \mathcal{P}_i(v_1, \dots, v_k - \omega, \dots, v_N), \quad \forall i \in 1..N,$$

$$\hat{\mathcal{P}}'_i(v_1, \dots, v_k, \dots, v_N) = \hat{\mathcal{P}}_i(v_1, \dots, v_k + \omega, \dots, v_N), \quad \forall i \in 1..N.$$

Equation (29) for $\hat{\mathcal{P}}^* = \mathcal{P}'$ takes the following form:

$$\begin{aligned} & \sum_{i=1, i \neq k}^N \int_{\mathbf{v}} v_i \hat{\mathcal{P}}_i(\mathbf{v}) \hat{G}(d\mathbf{v}) + \int_{\mathbf{v}} v_k \hat{\mathcal{P}}_k(\mathbf{v}) \hat{G}(d\mathbf{v}) - \lambda \kappa(\hat{\mathcal{P}}, \hat{G}) \geq \\ & \geq \sum_{i=1, i \neq k}^N \int_{\mathbf{v}} v_i \mathcal{P}'_i(\mathbf{v}) \hat{G}(d\mathbf{v}) + \int_{\mathbf{v}} v_k \mathcal{P}'_k(\mathbf{v}) \hat{G}(d\mathbf{v}) - \lambda \kappa(\mathcal{P}', \hat{G}). \end{aligned} \quad (30)$$

Now, we perform a transformation of coordinates: $(v_1, \dots, v_k, \dots, v_N) \rightarrow (u_1, \dots, u_k + \omega, \dots, u_N)$, which allows us to substitute G for \hat{G} , \mathcal{P} for \mathcal{P}' , and $\hat{\mathcal{P}}'$ for $\hat{\mathcal{P}}$.

$$\begin{aligned} & \sum_{i=1, i \neq k}^N \int_{\mathbf{u}} u_i \hat{\mathcal{P}}'_i(\mathbf{u}) G(d\mathbf{u}) + \int_{\mathbf{u}} (u_k + \omega) \hat{\mathcal{P}}'_k(\mathbf{u}) G(d\mathbf{u}) - \lambda \kappa(\hat{\mathcal{P}}', G) \geq \\ & \geq \sum_{i=1, i \neq k}^N \int_{\mathbf{u}} u_i \mathcal{P}_i(\mathbf{u}) G(d\mathbf{u}) + \int_{\mathbf{u}} (u_k + \omega) \mathcal{P}_k(\mathbf{u}) G(d\mathbf{u}) - \lambda \kappa(\mathcal{P}, G). \end{aligned} \quad (31)$$

Finally, we write the second term on the left hand side in the following form

$$\int_{\mathbf{u}} (u_k + \omega) \hat{\mathcal{P}}'_k(\mathbf{u}) G(d\mathbf{u}) = \int_{\mathbf{u}} u_k \hat{\mathcal{P}}'_k(\mathbf{u}) G(d\mathbf{u}) + \omega \int_{\mathbf{u}} \hat{\mathcal{P}}'_k(\mathbf{u}) G(d\mathbf{u}). \quad (32)$$

Equation (31) then states:

$$U(\hat{\mathcal{P}}', G) \geq U(\mathcal{P}, G) + (\mathcal{P}_k^0 - \hat{\mathcal{P}}_k^0) \omega. \quad (33)$$

Using the assumption, (25), we show the following

$$U(\hat{\mathcal{P}}', G) > U(\mathcal{P}, G), \quad (34)$$

which is a contradiction to equation (28). □

B.2 Problem 1

Proof of Theorem 2: The solution to the DM's problem is unique. This is due to Lemma 3 together with Assumption 1 in Appendix A.

The DM 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 (7), we arrive at the result. \square

B.3 Problem 3

Proof of Theorem 3: Let us consider a problem with $N + 1$ options, where the options N and $N + 1$ are duplicates. Let $\{\hat{\mathcal{P}}_i^0(\mathbf{u})\}_{i=1}^{N+1}$ be the unconditional probabilities in the solution to this problem. Since u_N and u_{N+1} are almost surely equal, then we can substitute u_N for u_{N+1} in the first order condition (7) to arrive at:

$$\hat{\mathcal{P}}_i(\mathbf{u}) = \frac{\hat{\mathcal{P}}_i^0 e^{u_i/\lambda}}{\sum_{j=1}^{N-1} \hat{\mathcal{P}}_j^0 e^{u_j/\lambda} + (\hat{\mathcal{P}}_N^0 + \hat{\mathcal{P}}_{N+1}^0) e^{u_N/\lambda}} \quad a.s., \forall i < N \quad (35)$$

$$\hat{\mathcal{P}}_N(\mathbf{u}) + \hat{\mathcal{P}}_{N+1}^0(\mathbf{u}) = \frac{(\hat{\mathcal{P}}_N^0 + \hat{\mathcal{P}}_{N+1}^0) e^{u_i/\lambda}}{\sum_{j=1}^{N-1} \hat{\mathcal{P}}_j^0 e^{u_j/\lambda} + (\hat{\mathcal{P}}_N^0 + \hat{\mathcal{P}}_{N+1}^0) e^{u_N/\lambda}} \quad a.s. \quad (36)$$

Therefore, the right hand sides do not change when only $\hat{\mathcal{P}}_N^0$ and $\hat{\mathcal{P}}_{N+1}^0$ change if their sum stays constant. Inspecting (35)-(36), we see that any such strategy produces the same expected value as the original one. Moreover, the amount of processed information is also the same for both strategies. To show this we use (7) to rewrite (4) as:¹⁷

$$\kappa = \int \sum_{i=1}^{N+1} \hat{\mathcal{P}}_i(\mathbf{u}) \log \frac{\hat{\mathcal{P}}_i(\mathbf{u})}{\hat{\mathcal{P}}_i^0} G(d\mathbf{u}) = \int \sum_{i=1}^{N+1} \hat{\mathcal{P}}_i(\mathbf{u}) \log \frac{e^{u_i/\lambda}}{\sum_{j=1}^{N-1} \hat{\mathcal{P}}_j^0 e^{u_j/\lambda} + (\hat{\mathcal{P}}_N^0 + \hat{\mathcal{P}}_{N+1}^0) e^{u_N/\lambda}} G(d\mathbf{u}). \quad (37)$$

Therefore, the achieved objective in (5) is the same for any such strategy as for the original strategy, and all of them solve the DM's problem.

¹⁷Here we use the fact that the mutual information between random variables X and Y can be expressed as $E_{p(x,y)} \left[\log \frac{p(x,y)}{p(x)p(y)} \right]$. See Cover and Thomas (2006, p. 20).

Finally, even the corresponding strategy with $\hat{\mathcal{P}}_{N+1}^0 = 0$ is a solution. Moreover, this implies that the remaining $\{\hat{\mathcal{P}}_i^0\}_{i=1}^N$ is the solution to the problem without the duplicate option $N + 1$, which completes the proof. \square

C Additional proofs and solutions

C.1 Problem 2

To solve the problem, we must find \mathcal{P}_1^0 , while $\mathcal{P}_2^0 = 1 - \mathcal{P}_1^0$. These probabilities must be internally consistent in that $\mathcal{P}_i^0 = \int_{\mathbf{v}} \mathcal{P}_i(\mathbf{v})G(d\mathbf{v})$, where $\mathcal{P}_i(\mathbf{v})$ is given by equation (7). Dividing each side of this condition by \mathcal{P}_i^0 yields:

$$1 = \frac{g_0}{\mathcal{P}_1^0 + \mathcal{P}_2^0 e^{\frac{R}{\lambda}}} + \frac{(1 - g_0)e^{\frac{1}{\lambda}}}{\mathcal{P}_1^0 e^{\frac{1}{\lambda}} + \mathcal{P}_2^0 e^{\frac{R}{\lambda}}} \quad \text{if } \mathcal{P}_1^0 > 0, \quad (38)$$

$$1 = \frac{g_0 e^{\frac{R}{\lambda}}}{\mathcal{P}_1^0 + \mathcal{P}_2^0 e^{\frac{R}{\lambda}}} + \frac{(1 - g_0)e^{\frac{R}{\lambda}}}{\mathcal{P}_1^0 e^{\frac{1}{\lambda}} + \mathcal{P}_2^0 e^{\frac{R}{\lambda}}} \quad \text{if } \mathcal{P}_2^0 > 0. \quad (39)$$

There are three solutions to this system,

$$\begin{aligned} \mathcal{P}_1^0 &\in \left\{ 0, 1, -\frac{e^{\frac{R}{\lambda}} \left(-e^{\frac{1}{\lambda}} + e^{\frac{R}{\lambda}} - g_0 + g_0 e^{\frac{1}{\lambda}} \right)}{\left(e^{\frac{1}{\lambda}} - e^{\frac{R}{\lambda}} \right) \left(-1 + e^{\frac{R}{\lambda}} \right)} \right\} \\ \mathcal{P}_2^0 &= 1 - \mathcal{P}_1^0. \end{aligned} \quad (40)$$

Now, we make an argument using the solution's uniqueness to deduce the true solution to the DM's problem. The first solution to the system, $\mathcal{P}_1^0 = 0$, corresponds to the case when the DM chooses option 2 without processing any information. The realized value is then R with certainty. The second solution, $\mathcal{P}_1^0 = 1$, results in the *a priori* selection of option 1 so the expected value equals $(1 - g_0)$. The third solution describes the case when the DM chooses to process a positive amount of information.

Problem 2 satisfies Assumption 2 as there are just two options and they do not take the same values with probability one. Therefore, Lemma 3 establishes that the solution to the

DM's optimization problem must be unique.

Since the expected utility is a continuous function of $\mathcal{P}_1^0, R, \lambda$ and g_0 , then the optimal \mathcal{P}_1^0 must be a continuous function of the parameters. Otherwise, there would be at least two solutions at the point of discontinuity of \mathcal{P}_1^0 . We also know that, when no information is processed, option 1 generates higher expected utility than option 2 for $(1 - g_0) > R$, and vice versa. So for some configurations of parameters $\mathcal{P}_1^0 = 0$ is the solution and for some configurations of parameters $\mathcal{P}_1^0 = 1$ is the solution. Therefore, the solution to the DM's problem has to include the non-constant branch, the third solution. To summarize this, the only possible solution to the DM's optimization problem is

$$\mathcal{P}_1^0 = \max \left(0, \min \left(1, -\frac{e^{\frac{R}{\lambda}} \left(-e^{\frac{1}{\lambda}} + e^{\frac{R}{\lambda}} - g_0 + g_0 e^{\frac{1}{\lambda}} \right)}{\left(e^{\frac{1}{\lambda}} - e^{\frac{R}{\lambda}} \right) \left(-1 + e^{\frac{R}{\lambda}} \right)} \right) \right). \quad (41)$$

C.2 Problem 4

To find the solution to Problem 4 we must solve for $\{\mathcal{P}_r^0, \mathcal{P}_b^0, \mathcal{P}_t^0\}$. The normalization condition $\mathcal{P}_r^0 = \int_{\mathbf{v}} \mathcal{P}_r(\mathbf{v})G(d\mathbf{v})$ yields:

$$1 = \frac{\frac{1}{4}(1+\rho)}{\mathcal{P}_r^0 + \mathcal{P}_b^0 + (1 - \mathcal{P}_r^0 - \mathcal{P}_b^0)e^{1/2\lambda}} + \frac{\frac{1}{4}(1-\rho)e^{1/\lambda}}{\mathcal{P}_r^0 e^{1/\lambda} + \mathcal{P}_b^0 + (1 - \mathcal{P}_r^0 - \mathcal{P}_b^0)e^{1/2\lambda}} \\ + \frac{\frac{1}{4}(1-\rho)}{\mathcal{P}_r^0 + \mathcal{P}_b^0 e^{1/\lambda} + (1 - \mathcal{P}_r^0 - \mathcal{P}_b^0)e^{1/2\lambda}} + \frac{\frac{1}{4}(1+\rho)e^{1/\lambda}}{\mathcal{P}_r^0 e^{1/\lambda} + \mathcal{P}_b^0 e^{1/\lambda} + (1 - \mathcal{P}_r^0 - \mathcal{P}_b^0)e^{1/2\lambda}} \quad (42)$$

Due to the symmetry between the buses, we know $\mathcal{P}_r^0 = \mathcal{P}_b^0$. This makes the problem one equation with one unknown, \mathcal{P}_r^0 . The problem can be solved analytically using the same arguments as in Appendix C.1. The resulting analytical expression is:

$$\mathcal{P}_r^0 = \max \left(0, \min \left(0.5, \frac{\begin{pmatrix} e^{\frac{1}{2\lambda}} - 8e^{\frac{1}{\lambda}} + 14e^{\frac{3}{2\lambda}} - 8e^{2/\lambda} + e^{\frac{5}{2\lambda}} \\ + \frac{1}{2}e^{\frac{1}{2\lambda}}(1-\rho) - e^{\frac{3}{2\lambda}}(1-\rho) + \frac{1}{2}e^{\frac{5}{2\lambda}}(1-\rho) \\ + e^{\frac{1}{2\lambda}}(-1 + e^{\frac{1}{\lambda}})x \end{pmatrix}}{2 \left(4e^{\frac{1}{2\lambda}} - 16e^{\frac{1}{\lambda}} + 24e^{\frac{3}{2\lambda}} - 16e^{2/\lambda} + 4e^{\frac{5}{2\lambda}} \right)} \right) \right),$$

where

$$x = \sqrt{\begin{pmatrix} 2 - 2e^{\frac{1}{\lambda}} + e^{2/\lambda} - 8e^{\frac{1}{2\lambda}}(1-\rho) + 14e^{\frac{1}{\lambda}}(1-\rho) \\ - 8e^{\frac{3}{2\lambda}}(1-\rho) + e^{2/\lambda}(1-\rho) + \frac{1}{4}(1-\rho)^2 \\ - \frac{1}{2}e^{\frac{1}{\lambda}}(1-\rho)^2 + \frac{1}{4}e^{2/\lambda}(1-\rho)^2 - \rho \end{pmatrix}}.$$

C.3 Inconsistency with a random utility model

This appendix establishes that the behavior of the rationally inattentive agent is not consistent with a random utility model. The argument is based on the counterexample described in section 4.4. Let Problem A refer to the choice among options 1 and 2 and Problem B refer to the choice among all three options. For simplicity, $\mathcal{P}_i(s)$ denotes the probability of selecting option i conditional on the state s , and $g(s)$ is the prior probability of state s .

Lemma 4. *For all $\epsilon > 0$ there exists Y s.t. the DM's strategy in Problem B satisfies*

$$\mathcal{P}_3(1) > 1 - \epsilon, \quad \mathcal{P}_3(2) < \epsilon.$$

Proof: For $Y > 1$, an increase of $\mathcal{P}_3(1)$ (decrease of $\mathcal{P}_3(2)$) and the corresponding relocation of the choice probabilities from (to) other options increases the agent's expected payoff. The resulting marginal increase of the expected payoff is larger than $(Y - 1) \min(g(1), g(2))$. Selecting Y allows us to make the marginal increase arbitrarily large and therefore the marginal value of information arbitrarily large.

On the other hand, with λ being finite, the marginal change in the cost of information is also finite as long as the varied conditional probabilities are bounded away from zero. See equation (4), the derivative of entropy with respect to $\mathcal{P}_i(s)$ is finite at all $\mathcal{P}_i(s) > 0$. Therefore, for any ϵ there exists high enough Y such that it is optimal to relocate probabilities from options 1 and 2 unless $\mathcal{P}_3(1) > 1 - \epsilon$, and to options 1 and 2 unless $\mathcal{P}_3(2) < \epsilon$. \square

Proof of proposition 3: we will show that there exist $g(1) \in (0, 1)$ and $Y > 0$ such that option 1 has zero probability of being selected in Problem A, while the probability is positive in both states in Problem B. Let us start with Problem A. According to Proposition 2, there exists a sufficiently high $g(1) \in (0, 1)$, call it g_M , such that the DM processes no information and $\mathcal{P}_1(1) = \mathcal{P}_1(2) = 0$. We will show that for $g(1) = g_M$ there exists a high enough Y , such the choice probabilities of option 1 are positive in Problem B.

Let $\mathcal{P} = \{\mathcal{P}_i(s)\}_{i=1,s=1}^{3,2}$ be the solution to Problem B. We now show that the optimal choice probabilities of options 1 and 2, $\{\mathcal{P}_i(s)\}_{i=1,s=1}^{2,2}$, solve a version of Problem A with modified prior probabilities. The objective function for Problem B is

$$\begin{aligned} \max_{\{\mathcal{P}_i(s)\}_{i=1,s=1}^{3,2}} & \sum_{i=1}^3 \sum_{s=1}^2 v_i(s) \mathcal{P}_i(s) g(s) \\ & - \lambda \left[- \sum_{s=1}^2 g(s) \log g(s) + \sum_{i=1}^3 \sum_{s=1}^2 \mathcal{P}_i(s) g(s) \log \frac{\mathcal{P}_i(s) g(s)}{\sum_{s'} \mathcal{P}_i(s') g(s')} \right], \end{aligned} \quad (43)$$

where we have written the information cost as $H(s) - E[H(s|i)]$.¹⁸ If $\mathcal{P}_3(1)$ and $\mathcal{P}_3(2)$ are the conditional probabilities of the solution to Problem B, the remaining conditional probabilities solve the following maximization problem.

$$\max_{\{\mathcal{P}_i(s)\}_{i=1,s=1}^{2,2}} \sum_{i=1}^2 \sum_{s=1}^2 v_i(s) \mathcal{P}_i(s) g(s) - \lambda \left[\sum_{i=1}^2 \sum_{s=1}^2 \mathcal{P}_i(s) g(s) \log \frac{\mathcal{P}_i(s) g(s)}{\sum_{s'} \mathcal{P}_i(s') g(s')} \right], \quad (44)$$

subject to $\mathcal{P}_1(s) + \mathcal{P}_2(s) = 1 - \mathcal{P}_3(s)$, $\forall s$. Equation (44) is generated from (43) by omitting the

¹⁸Recall that $H(Y|X) = - \sum_{x \in X} \sum_{y \in Y} p(x, y) \log p(y|x)$ (Cover and Thomas, 2006, p. 17).

terms independent of $\{\mathcal{P}_i(s)\}_{i=1,s=1}^{2,2}$. Now, we make the following substitution of variables.

$$\mathcal{R}_i(s) = \mathcal{P}_i(s) / \left(1 - \mathcal{P}_3(s)\right) \quad (45)$$

$$\hat{g}(s) = Kg(s) \left(1 - \mathcal{P}_3(s)\right) \quad (46)$$

$$1/K = \sum_{s=1}^2 g(s) \left(1 - \mathcal{P}_3(s)\right). \quad (47)$$

where K , which is given by (47), is the normalization constant that makes the new prior, $\hat{g}(s)$, sum up to 1.

The maximization problem (44) now takes the form:

$$\max_{\{\mathcal{R}_i(s)\}_{i=1,s=1}^{1,2}} \sum_{i=1}^2 \sum_{s=1}^2 v_i(s) \mathcal{R}_i(s) \hat{g}(s) - \lambda \sum_{i=1}^2 \sum_{s=1}^2 \mathcal{R}_i(s) \hat{g}(s) \log \frac{\mathcal{R}_i(s) \hat{g}(s)}{\sum_{s'} \mathcal{R}_i(s') \hat{g}(s')}, \quad (48)$$

subject to

$$\mathcal{R}_1(s) + \mathcal{R}_2(s) = 1 \quad \forall s. \quad (49)$$

The objective function of this problem is equivalent to (44) up to a factor of K , which is a positive constant. The optimization problem (48) subject to (49) is equivalent to Problem A with the prior modified to from $g(s)$ to $\hat{g}(s)$, let us call it Problem C.¹⁹

According to Proposition 2, there exists $\hat{g}_m \in (0, 1)$ such that the DM always selects option 1 in Problem C for all $\hat{g}(1) \leq \hat{g}_m$. From equations (46) and (47) we see that for any $\hat{g}_m > 0$ and $g(1), g(2) \in (0, 1)$ there exists $\epsilon > 0$ such that if $\mathcal{P}_3(1) > 1 - \epsilon$ and $\mathcal{P}_3(2) < \epsilon$, then $\hat{g}(1) < \hat{g}_m$.²⁰ Moreover, Lemma 4 states that for any such $\epsilon > 0$ there exists Y such that $\mathcal{P}_3(1) > 1 - \epsilon$ and $\mathcal{P}_3(2) < \epsilon$. Therefore there is a Y such that in Problem C, option 1 is selected with positive probability in both states, which also implies it is selected with positive probabilities in Problem B, see equation (45). \square

¹⁹To see the equivalence to Problem A, observe that this objective function has the same form as (43) except for a) the constant corresponding to $H(s)$ and b) we only sum over $i = 1, 2$.

²⁰ $\hat{g}(1) = \frac{g(1)(1-\mathcal{P}_3(1))}{\sum_s g(s)(1-\mathcal{P}_3(s))} < \frac{g(1)(1-\mathcal{P}_3(1))}{g(2)(1-\mathcal{P}_3(2))} < \frac{g(1)\epsilon}{g(2)(1-\epsilon)}$.