

NBER WORKING PAPER SERIES

INFORMATION COSTS AND SEQUENTIAL INFORMATION SAMPLING

Benjamin Hébert
Michael Woodford

Working Paper 25316
<http://www.nber.org/papers/w25316>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2018

The authors would like to thank Mark Dean, Sebastian Di Tella, Mira Frick, Matthew Gentzkow, Emir Kamenica, Divya Kirti, Jacob Leshno, Stephen Morris, Pietro Ortoleva, José Scheinkman, Ilya Segal, Ran Shorrer, Miguel Villas-Boas, Ming Yang, and participants at the Cowles Theory conference, 16th SAET Conference, Barcelona GSE Summer Conference on Stochastic Choice, Stanford GSB research lunch, the 2018 ASSA meetings, and UC Berkeley Theory Seminar for helpful discussions on this topic, and the NSF for research support. We would particularly like to thank Doron Ravid and Philipp Strack for discussing an earlier version of the paper. We would also like to thank the editor and anonymous referees for comments that significantly improved the paper. Portions of this paper circulated previously as the working paper “Rational Inattention with Sequential Information Sampling,” which has been split in two, and appeared in Benjamin Hébert’s Ph.D. dissertation at Harvard University. All remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Benjamin Hébert and Michael Woodford. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Information Costs and Sequential Information Sampling
Benjamin Hébert and Michael Woodford
NBER Working Paper No. 25316
November 2018
JEL No. D83,E70

ABSTRACT

We propose a new approach to modeling the cost of information structures in rational inattention problems, the "neighborhood-based" cost functions. These cost functions have two properties that we view as desirable: they summarize the results of a sequential evidence accumulation problem, and they capture notions of "perceptual distance." The first of these properties is connected to an extensive literature in psychology and neuroscience, and the second ensures that neighborhood-based cost functions, unlike mutual information, make accurate predictions about behavior in perceptual experiments. We compare the implications of our neighborhood-based cost functions with those of a mutual-information cost function in a series of applications: security design, global games, modeling perceptual judgments, and a linear-quadratic-Gaussian tracking problem.

Benjamin Hébert
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
bhebert@stanford.edu

Michael Woodford
Department of Economics
Columbia University
420 W. 118th Street
New York, NY 10027
and NBER
mw2230@columbia.edu

A data appendix is available at <http://www.nber.org/data-appendix/w25316>

1 Introduction

In models of rational inattention (proposed by Christopher Sims and surveyed in Sims (2010)), a decision maker (DM) chooses her action based on a signal that provides only an imperfect indication of the true state. The information structure that generates this signal is optimal, in the sense of allowing the best possible state-contingent action choice, net of a cost of information. In Sims' theory, the cost of any information structure is proportional to the Shannon mutual information between the true state of the world and the signals generated by that information structure.

It is not obvious, though, that the theorems that justify the use of mutual information in communications engineering (Cover and Thomas (2012)) provide any warrant for using it as a cost function in a theory of attention allocation, either in the case of economic decisions or that of perceptual judgments.¹ Moreover, the mutual-information cost function has implications that are unappealing on their face, and that seem inconsistent with evidence on the nature of sensory processing, as discussed, for example, in Woodford (2012), Caplin and Dean (2013), Dewan and Neligh (2017), and Caplin et al. (2018b).

We propose an alternative family of information costs, which we call “neighborhood-based” cost functions. These information costs have two particular properties (in addition to the standard ones described in, e.g., De Oliveira et al. (2017)) that we view as desirable. First, they can be viewed as summarizing the results of a process of sequential evidence accumulation, in which each successive increment to the cumulatively available evidence is only very minimally informative. Second, these information costs can capture the idea that certain pairs of states are easy to distinguish, whereas others are difficult to distinguish. This second property allows the neighborhood-based cost functions avoid some of the prob-

¹As explained in Cover and Thomas (2012), these theorems rely upon the possibility of “block coding” of a large number of independent instances of a given type of message, that can be jointly transmitted before any of the messages have to be decoded by the recipient. In our situation, an action must be taken in an individual decision problem, without waiting to learn about a large number of problems of the same form.

lematic implications of mutual information. The two properties are connected by an object we call the “information-cost matrix function,” which encodes the difficulty of distinguishing between pairs of states and summarizes the cost of a small amount of information in the sequential evidence accumulation problem.

The neighborhood-based costs functions differ from mutual information because mutual information imposes a type of symmetry across different states of nature, so that it is equally difficult to distinguish between any two states that are equally probable *ex ante*. This implies that under an optimal information structure, actions differ across states only to the extent that the associated payoffs differ across those states, and action probabilities jump discontinuously when payoffs jump. An extensive experimental literature in psychophysics finds that subjects’ probabilities of making perceptual judgments (the action) vary continuously with changes in the stimulus magnitude along some dimension (the state), even when subjects are rewarded based on whether the magnitude is greater or smaller than some threshold (generating a discrete jump in payoffs).² Such behavior can be optimal only if it is costly to have an information structure that generates very different signals in similar states, while making it less costly to distinguish states that are dissimilar. In other words, the information cost must capture some notion of “perceptual distance.”

Motivated by these issues, we consider the properties that a plausible cost function should satisfy. As discussed in Fehr and Rangel (2011) and Woodford (2014), a large literature in psychology and neuroscience has argued that data on both the frequency of perceptual errors and the frequency distribution of response times can be explained by models of sequential sampling. More recently, some authors have proposed that data on stochastic choice and response time in economic contexts can be similarly modeled.³ Consequently, one property we desire in a cost function is that it should summarize the results

²See further discussion in section 5.2.

³Additional recent examples include Krajbich et al. (2014) and Clithero (2018). Shadlen and Shohamy (2016) provide a neural-process interpretation of sequential-sampling models of choice.

of a sequential evidence accumulation process. In this paper, we begin with a continuous-time model of sequential evidence accumulation, and then show that the resulting state-contingent choice probabilities are identical to those of a static rational inattention model with a uniformly posterior-separable cost function.⁴ We derive this continuous time diffusion model from a discrete time model in Hébert and Woodford (2018).⁵

The use of diffusion processes to model the evolution of an internal belief state has been popular in the aforementioned literature from mathematical psychology and computational neuroscience, and in the economic applications reviewed by Fehr and Rangel (2011). In that literature, however, it is common to take as exogenous both the dynamics of the belief state as a function of the true decision situation, and the criterion used to decide when to stop deliberating and make a decision. Our goal instead is to derive both the nature of instantaneous evidence accumulation (and hence the dynamics of the resulting belief state) and the stopping rule (as well as the rule that determines the action to be taken) from an optimization principle, under an explicit model of the cost of information gathering.

Our paper is not the first that seeks to derive at least some features of such models from a theory of optimal information sampling. Moscarini and Smith (2001) consider both the optimal intensity of information sampling per unit of time and the optimal stopping problem, when the only possible kind of information is given by the sample path of a Brownian motion with a drift that depends on an unknown state.⁶ Fudenberg et al. (forthcoming) consider a variant of this problem with a continuum of possible states, and an exogenously

⁴For more on this class of cost functions, see Caplin et al. (2018b). Morris and Strack (2017) provide a related foundation for this class, in the special case in which there are only two possible states and signals are exogenous. Note, however, that with only two states, because beliefs must diffuse on a line, there is little distinction between “endogenous” and “exogenous” signals.

⁵In section C.1 of the Technical Appendix, we consider an extension of our results to a continuous-time evidence accumulation process with Poisson jumps (also derived in Hébert and Woodford (2018)), as in the models of optimal evidence accumulation proposed by Che and Mierendorff (2017) and Zhong (2018). For a discussion of the differences between these approaches, see Hébert and Woodford (2018).

⁶Moscarini and Smith (2001) allow the instantaneous variance of the observation process to be freely chosen (subject to a cost), but this is equivalent to changing how much of the sample path of a given Brownian motion can be observed by the DM within a given amount of clock time.

fixed sampling intensity.⁷ Woodford (2014) instead takes as given a stopping rule, but allows a flexible choice of the information sampling process. Our approach differs from these earlier efforts in seeking to endogenize *both* the nature of the information that is sampled at each stage of the evidence accumulation process and the stopping rule that determines how much evidence is collected before a decision is made. We also consider decision problems with an arbitrary number of choice alternatives, as opposed to the binary choice problems of Fudenberg et al. (forthcoming) and Woodford (2014).

The key “parameter” of our diffusion model is a matrix-valued function that describes the local cost of information acquisition. If the DM chooses to have her beliefs diffuse with higher variance in some dimension, she is gathering information in this dimension, and the information-cost matrix function describes the cost of this information. It encodes, on its diagonal, how difficult each state is to learn about, and on its off-diagonal, how difficult it is to discriminate between states. Mutual information generates problematic predictions because its corresponding information-cost matrix function has a kind of symmetry that implies that equally likely states are equally difficult to discriminate.

The information-cost matrix function provides a bridge between the first property we consider desirable (equivalence to a sequential problem) and the second property (capturing perceptual distance). Given an information cost function in the sequential evidence accumulation problem, there is a static rational inattention problem with a particular uniformly posterior separable cost function that generates the equivalent joint distribution of actions and states. Moreover, the comparative statics of this joint distribution with respect to changes in payoffs are governed by the information cost matrix function. Intuitively, if the information cost matrix function makes it very costly to discriminate between some pair of states, the DM will not do so even if her payoff jumps across those states. As a result, there are uniformly posterior separable cost functions that can both summarize sequential

⁷See also Tajima et al. (2016) for analysis of a related class of models.

evidence accumulation and capture the idea of perceptual distance.

We introduce a family of such cost functions, the neighborhood-based cost functions. The idea of this class of information-cost specifications is that information structures are more costly the greater the extent to which they allow intrinsically similar states of the world (states that share a “neighborhood”) to be discriminated; the dependence on a concept of intrinsic similarity between states (the “neighborhood structure”) distinguishes these cost functions from the mutual information cost function. We show that using information costs in this family can explain the continuous variation of response frequencies in the perceptual experiments mentioned previously. Dean and Neligh (2018) study neighborhood-based cost functions in an experimental setting, and find that these costs fit observed behavior better than several other alternatives, including mutual information.

We also specialize this family to a particularly useful case, in which the states can be ordered on a line. Throughout the paper, we use as a running example the case of a potential buyer of a security whose payoff depends on the value of some assets (an example based on Yang (2017)). In this case, it is natural to suppose that the relevant states of the world are the asset values, and that it may be difficult for the DM to discriminate between nearby asset values even as the DM is more easily able to acquire information about whether the asset values will be very high or very low. We extend our analysis of this case to a continuum of states (in the rest of the paper, we use a discrete state space) and show that the limit of the neighborhood-based cost function for this neighborhood structure is the average Fisher information. This is the average value over the state space of a local measure of the discriminability of nearby states. Like mutual information, this measure is uniquely defined up to a scale parameter, and it can be used instead of mutual information in almost any context in which the states can be ordered on a line or a circle.

After we introduce the neighborhood-based cost functions, and show that they satisfy the two properties of cost functions that we view as desirable, we discuss four applications

that illustrate how they are different from or similar to mutual information. We study perceptual experiments, global games (building on Morris and Yang (2016)), security design (building on Yang (2017)), and a linear-quadratic-Gaussian setting (as in Sims (2010)). In the first three of these, the neighborhood-based cost functions generate different predictions than mutual information. In the popular linear-quadratic-Gaussian case, reassuringly, mutual information and the average Fisher information generate identical predictions.

Section 2 presents our continuous-time model of sequential evidence accumulation, and introduces the information cost matrix function as a way of parameterizing flow information costs. Section 3 proves the equivalence of sequential evidence accumulation and static rational inattention problems with uniformly posterior-separable costs. In section 4, we introduce the neighborhood-based cost functions. We apply the neighborhood-based cost functions in a series of applications in section 5. In section 6 we conclude.

2 Continuous-Time Sequential Evidence Accumulation

We begin by introducing our continuous-time model of sequential evidence accumulation. We derive this model from a discrete-time dynamic evidence accumulation problem in Hébert and Woodford (2018). Our derivation depends on a few key assumptions, whose import we discuss below. During the setup of the model, we will use as an example our security design application (Section 5.1), which is based on Yang (2017).

Let $x \in X$ be the underlying state of the nature, and $a \in A$ be the action taken by the decision maker (DM). For simplicity, A and X are finite sets, and the number of states is weakly larger than the number of actions, $|X| \geq |A|$. The DM's utility from taking action a in state x at time t is $u_{a,x} - \kappa t$. The parameter $\kappa > 0$ governs the penalty for delaying making a decision; the DM does not discount the future.

The DM does not initially know the state $x \in X$, but can learn about which states are

more or less likely. At each time t , the DM holds beliefs $q_t \in \mathcal{P}(X)$, where $\mathcal{P}(X) \subset \mathbb{R}^{|X|}$ denotes the probability simplex over X . That is, q_t is a vector of length $|X|$, whose elements, denoted $q_{x,t}$, are the probability, under the DM's beliefs at time t , of state x . Time begins at $t = 0$, when the DM holds prior beliefs q_0 . At each moment in time, the DM faces two decisions: whether to gather information about the state $x \in X$, and whether to stop and make a decision. When stopping with beliefs q_τ at time τ , the DM will choose a to maximize $u_a^T \cdot q_\tau$, where u_a is the vector of utilities associated with action a . Define

$$\hat{u}(q_\tau) = \max_{a \in A} u_a^T \cdot q_\tau,$$

and note that $\hat{u}(q_\tau) - \kappa\tau$ is the payoff if the DM stops with beliefs q_τ at time τ and then chooses the optimal action given those stopping beliefs.

Example. Suppose the DM is considering buying a security whose payoff is a function of the value of some assets. In this case, X is a set of possible values for the assets, and the actions are to either accept (L , “like”) or reject (R) the offer, $A = \{L, R\}$. The utility of rejecting the offer is normalized to zero ($u_{R,x} = 0$), and the utility of accepting the offer is $u_{L,x} = s_x - K$, where s_x is the security payoff and K is the price. The stopping payoff $\hat{u}(\cdot)$ involves deciding, under the current beliefs, whether to accept or reject: $\hat{u}(q_\tau) = \max\{q_\tau^T s - K, 0\}$, where s is the vector of security payoffs.

When the DM gathers information, she chooses the variance-covariance matrix of possible changes in her beliefs, subject to certain constraints. The DM's beliefs evolve as

$$dq_{x,t} = q_{x,t} \sigma_{x,t} \cdot dB_t, \tag{1}$$

where dB_t is an $|X|$ -dimensional Brownian motion,⁸ σ_t is a matrix that can be chosen by

⁸The $|X|$ -th dimension is redundant since beliefs stay in the simplex; we keep it for notational convenience.

the DM, and $\sigma_{x,t}$ is a particular row of that matrix. This process is derived from Bayesian updating, and imposes two restrictions on the DM's beliefs. First, beliefs are martingales, and second, $dq_{x,t} = 0$ if $q_{x,t} = 0$, meaning that if the DM believes a particular state is not possible, she will never come to believe that state is possible at some point in the future.

The DM's choice of σ_t is subject to restrictions — a trivial one to ensure that the beliefs stay in the simplex, and an economic restriction that limits the amount of information the DM can acquire. The trivial restriction is that

$$\mathbf{1}^T \cdot dq_t = 0$$

always, where $\mathbf{1}$ is a vector of ones. This restriction is equivalent to requiring that $\sigma_t^T q_t = \vec{0}$. We will use $M(q_t)$ to denote the set of $|X| \times |X|$ matrices satisfying this condition.

The non-trivial restriction limits the quantity of information the DM can acquire:

$$\frac{1}{2} \text{tr}[\sigma_t^T k(q_t) \sigma_t] \leq \chi, \quad (2)$$

where $k(q_t)$ is an $|X| \times |X|$ dimensional matrix-valued function we will refer to as the “information-cost matrix function”, $\text{tr}[\cdot]$ is the trace, and χ is a positive constant that indexes the tightness of the constraint. We discuss this constraint, and the information-cost matrix function, in more detail below. Our derivation in Hébert and Woodford (2018) shows that the information-cost matrix function satisfies certain properties: for any q_t , $k(q_t)$ is symmetric and positive semi-definite, and its null space is the space of vectors that are constant for all $x \in X$ in the support of q_t . Because the information cost matrix function $k(q_t)$ is strictly positive-definite outside of this space, every matrix $\sigma_t \in M(q_t)$ that generates volatility in beliefs also generates a strictly positive value for $\text{tr}[\sigma_t^T k(q_t) \sigma_t]$.⁹

⁹Lemma 1 below proves (among other things) this fact.

Using her control of the volatility of her beliefs, and subject to the constraints imposed by the information-cost matrix function, our DM attempts to maximize her expected payoff. Her sequence problem can be written, given beliefs q_t at time t ,

$$V(q_t) = \sup_{\{\sigma_s \in M(q_s)\}, \tau \geq t} E_t[\hat{u}(q_\tau) - \kappa(\tau - t)],$$

where τ is the DM's endogenous stopping time, subject to the constraint (2).

Anywhere this value function is twice-differentiable and the DM does not choose to stop, the Hamilton-Jacobi-Bellman (HJB) equation associated with this problem is

$$\begin{aligned} \sup_{\sigma_t \in M(q_t)} \frac{1}{2} \text{tr}[\sigma_t^T \text{Diag}(q_t) V_{qq}(q_t) \text{Diag}(q_t) \sigma_t] &= \kappa, \\ \text{subject to } \frac{1}{2} \text{tr}[\sigma_t^T k(q_t) \sigma_t] &\leq \chi, \end{aligned} \quad (3)$$

where $\text{Diag}(q_t)$ is a diagonal matrix with the elements of q_t on its diagonal, and $V_{qq}(q_t)$ is the Hessian of $V(q)$ evaluated at $q = q_t$.¹⁰

The DM's optimal stopping rule is characterized by the standard value-matching and smooth-pasting conditions. Let $\Omega \subset \mathcal{P}(X)$ be the open subset of the simplex on which the DM continues to search for information, and let $\partial\Omega$ denote its boundary. For all $q \in \partial\Omega$, the value matching condition, $V(q) = \hat{u}(q)$, and smooth pasting condition, $V_q(q) = \hat{u}_q(q)$, will hold. Note, however, that the derivative $\hat{u}_q(q)$ does not exist everywhere — at beliefs where the DM is just indifferent between two actions with distinct state-contingent payoffs, the stopping payoff is non-differentiable.¹¹ However, it will never be optimal for the DM

¹⁰The function $V(q)$ is defined on the probability simplex $\mathcal{P}(X)$, but we find it convenient to extend it to the space of measures, $\mathbb{R}_+^{|X|}$, by assuming it is homogeneous of degree one. This assumption does not restrict the function's values on the simplex. Under this assumption, vectors in the tangent space are simply vectors in $\mathbb{R}^{|X|}$, which we express using the natural set of basis vectors corresponding to each element of X . The Hessian matrices appearing in equations such as (3) above and (10) below should be understood in this way.

¹¹We have yet to show that $V(q)$ is differentiable everywhere, but prove this as part of Theorem 1.

to stop at one of these indifference points.

Next, we provide some intuition for the volatility constraint and the information-cost matrix function. The constraint is a limit on the information the DM can acquire, because it limits the volatility of her beliefs. Our DM is a Bayesian, meaning that she can never expect to revise her beliefs in a particular direction — her beliefs must be a martingale; this is why there can be no drift term in equation (1). If she receives a mostly uninformative signal at a particular moment, her beliefs have a small amount of volatility at that moment. In contrast, if she receives an informative signal, her beliefs will be very volatile.

We derive the information constraint (2) from a model in which the DM can choose any information structure each time period, as in standard rational inattention models. One result of our derivation is that the DM can choose any volatility matrix σ_t . This is, in a sense, a familiar idea — Kamenica and Gentzkow (2011), for example, emphasize the idea of choosing a distribution of posteriors, subject to the constraint that the mean posterior is equal to the prior. Our DM appears to choose only the volatility, and not the higher cumulants of the distribution of posteriors, but this is because she finds it optimal to smooth her information gathering over time, and the instantaneous volatility is sufficient to characterize the resulting process for beliefs. This result permits both a relatively parsimonious specification of the information sampling strategies available to the DM, and a relatively parsimonious specification of possible forms for the information constraint.

Because we model the evolution of the DM’s beliefs as a diffusion process, our model resembles, e.g., Krajbich et al. (2014) and Fudenberg et al. (forthcoming). Unlike those authors, we endogenize the diffusion process through which additional information arrives while sampling continues. Additionally, our model emphasizes the “unconditional” dynamics of beliefs (that is, not conditional on any particular state being the true state), whereas the models discussed by those authors are described in terms of their “conditional” dynamics (that is, conditional on some particular state being the true state).

The information-cost matrix function $k(q_t)$ is more than simply a way of obtaining a single (scalar) measure of the “size” of the elements of σ_t . The relative size of different elements of the matrix also allows us to specify the degree to which it is more relatively costly to obtain certain kinds of information. In Hébert and Woodford (2018), we construct the matrix function $k(q)$ from the cost of receiving signals about the different states. The diagonal elements k_{xx} control the cost of receiving signals, conditional on the true state being x , that differ from the unconditional distribution of signals. The elements k_{xx} are always positive, and the larger they are, the more costly it is to distinguish the state x from the other states. The off-diagonal elements, $k_{xx'}$, control the cost of receiving signals conditional on state x that differ from the signals received conditional on x' . If $k_{xx'} < 0$, receiving similar signals conditional on x and x' reduces the overall information cost, and the magnitude of $k_{xx'}$ controls the size of this effect. For this reason, we think of $-k_{xx'}$ as a measure of how difficult the states x and x' are to discriminate.¹²

An example of an information-cost matrix function that satisfies our assumptions (and will be important for the discussion below) is the “inverse Fisher information matrix,”

$$k(q) = g^+(q) = \text{Diag}(q) - qq^T = \begin{bmatrix} q_1(1 - q_1) & -q_1q_2 & \dots & -q_1q_{|X|} \\ -q_1q_2 & q_2(1 - q_2) & \dots & -q_2q_{|X|} \\ \vdots & \vdots & \ddots & \vdots \\ -q_1q_{|X|} & -q_2q_{|X|} & \dots & q_{|X|}(1 - q_{|X|}) \end{bmatrix}. \quad (4)$$

In this case, the off-diagonal element $k_{xx'}(q)$ is equal to $-q(x)q(x')$ for any pair of states x, x' ; thus it depends only on the prior probabilities of the two states, and is otherwise the same regardless of the states selected. Consequently, all pairs of distinct states with identical probabilities are assumed to be equally easy or difficult to tell apart. While this

¹²If $k_{xx'} > 0$, receiving very different signals conditional on x and x' reduces the information cost. We cannot rule this out on theoretical grounds, but none of our examples feature positive off-diagonal elements.

kind of symmetry might seem appealing on a priori grounds for some applications, we view it as implausible for many cases of economic relevance.

Example. Continuing the example of a buyer considering a security, suppose the buyer's current beliefs q_t are uniformly distributed over the various asset values $x \in X$. If $k(q)$ is the inverse Fisher information matrix, the buyer finds it equally costly to discriminate between any pair of asset values x, x' , regardless of how close or far apart those asset values are.

In many applications, we have a notion of some pairs of states x, x' being closer or farther apart than others. In the case of payoffs, quantities, or other economic variables that can be summarized by a single number, we usually think that it is harder to sharply discriminate between values that are close together than values that are far apart. Perceptual experiments, in which subjects are asked to classify stimuli that differ from one another in intensity or magnitude along a single dimension, are another application with this feature.

An alternative information-cost matrix function, also satisfying our assumptions, is

$$k(q) = \begin{bmatrix} \frac{q_1 q_2}{q_1 + q_2} & -\frac{q_1 q_2}{q_1 + q_2} & 0 & \dots & 0 \\ -\frac{q_1 q_2}{q_1 + q_2} & \frac{q_1 q_2}{q_1 + q_2} + \frac{q_2 q_3}{q_2 + q_3} & -\frac{q_2 q_3}{q_2 + q_3} & \ddots & \vdots \\ 0 & -\frac{q_2 q_3}{q_2 + q_3} & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & \frac{q_{|X|-1} q_{|X|-2}}{q_{|X|-2} + q_{|X|-1}} + \frac{q_{|X|} q_{|X|-1}}{q_{|X|-1} + q_{|X|}} & -\frac{q_{|X|-1} q_{|X|}}{q_{|X|} + q_{|X|-1}} \\ 0 & \dots & 0 & -\frac{q_{|X|} q_{|X|-1}}{q_{|X|} + q_{|X|-1}} & \frac{q_{|X|-1} q_{|X|}}{q_{|X|} + q_{|X|-1}} \end{bmatrix}. \quad (5)$$

In this case, the only non-zero off-diagonal elements $k_{xx'}(q)$ are negative elements when x' directly follows x in the ordering of states (or vice versa). This form of matrix $k(q)$ implies that an information structure is costly only to the extent that there are pairs of “neighboring” states x, x' for which the distribution of signals conditional on those states are different. This information cost matrix function is closely related to the neighborhood-based cost functions we introduce in Section 4.

Example. Continuing the example of a buyer considering a security, if $k(q)$ is the “neighborhood-based” function described in equation (5) above, the buyer finds discriminating between adjacent asset values costly, and the total information cost depends on how rapidly the signals the buyer receives change as a function of the asset value.

Aside from its a priori appeal, this alternative information-cost matrix function has different implications for the behavior of the DM, which we describe in the next section.

For a large class of information-cost matrix functions $k(q_t)$, including both of these examples, we can solve the continuous time sequence problem described in this section, and show that the solution is equivalent to a certain static rational inattention problem. We present these results in the next section. The remainder of this section briefly discusses an additional restriction on the functions $k(q)$, the importance of linear time costs as opposed to exponential discounting, and our assumption that beliefs follow a diffusion process.

In Hébert and Woodford (2018), we derive this model from more primitive considerations, and our derivation shows that there exists a positive constant m such that $k(q) - mg^+(q)$ is positive semi-definite.¹³

We also derive a more general continuous time problem, involving a controlled jump-diffusion process for beliefs, that features both exponential (standard) discounting and linear time costs. The model presented in this paper corresponds to a special case, in which only linear time costs are present and the information cost technology exhibits what we call a “preference for gradual learning.” Under these assumptions, it is always weakly optimal for the DM to choose not to have jumps in her beliefs, meaning that her beliefs will follow a pure diffusion process. No-discounting generates a great deal of tractability (and decision times are often short relative to the rate of time preference), and this “preference for gradual learning” allows us to focus on models that are closely related to existing models in psychology, neuroscience, and economics.

¹³The example in (5) does not satisfy this for any $m > 0$, but is the limit of a sequence that does.

Other authors (e.g. Che and Mierendorff (2017), Zhong (2018)) have explored models in which beliefs are assumed to follow a jump process. In Hébert and Woodford (2018), we discuss conditions under which beliefs will follow a diffusion-like or jump-like process, with and without exponential discounting, and how our results relate to those papers. For robustness, in this paper, we show in Appendix Section C.1 that our equivalence result (Theorem 1 below) can be derived from an alternative model in which beliefs follow a pure jump process. Our subsequent results depend only on this theorem and thus are relevant to both diffusion-based and jump-based models. To preserve this robustness, we do not discuss topics like stopping times, which likely differ between diffusion-based and jump-based models, even when those models are both equivalent to the same static model.

3 Static and Dynamic Rational Inattention Problems

In most theories of rational inattention, including the classic formulations of Sims, only a single signal is collected for each decision that must be made. In a decision problem where an action is chosen once from a set of possibilities, the rational inattention problem is static; a signal is obtained (once) that depends on the state, an action is taken that depends on the signal, and that is all. The kind of dynamic optimization model proposed in the previous section seems quite different. Nonetheless, we establish below that in a broad class of cases, there is an equivalence between the information that is ultimately acquired in the dynamic model of the previous section and the information acquired in a static model of rational inattention, with a particular type of cost function. Thus, our dynamic model does not necessarily have different implications, on some dimensions, than a static rational inattention model; however, the dynamic optimization problem can provide a reason for interest in static information-cost functions of particular types.

We begin by explaining the form of a static rational inattention problem, and then de-

scribe our equivalence result. After introducing our equivalent result, we will discuss comparative statics, and relate those comparative statics to the information cost matrix function introduced in the previous section.

3.1 Static Models of Rational Inattention

Much of the notation from the previous section carries over to static models. We continue to use $x \in X$ as the underlying state of nature, and $u_{a,x}$ as the payoff from taking action $a \in A$ in state x . We continue to define $\hat{u}(q)$ as the payoff from taking the optimal action with beliefs q , and let q_0 denote the DM's initial beliefs, prior to gathering information.

In static rational inattention models, the DM chooses a “signal structure,” consisting of a signal alphabet S (a finite set) and a conditional probability, for each state x , of each signal, $p = \{p_x \in \mathcal{P}(S)\}_{x \in X}$. The signal structure p generates, under the prior beliefs q_0 , an unconditional probability of each signal, $\pi_s(p, q_0)$. After receiving a signal $s \in S$, the DM will hold beliefs $q_s(p, q_0)$, defined by Bayes' rule. Let $C(p, q_0; S) : \mathcal{P}(S)^{|X|} \times \mathcal{P}(X) \rightarrow \mathbb{R}$ be the cost of choosing a signal structure p and alphabet S , given initial prior q_0 . The standard static rational inattention problem, given the signal alphabet S ,¹⁴ is

$$\max_{\{p_x \in \mathcal{P}(S)\}_{x \in X}} \sum_{s \in S} \pi_s(p, q_0) \hat{u}(q_s(p, q_0)) - \theta C(p, q_0; S), \quad (6)$$

where $\theta > 0$ is a multiplicative factor that parameterizes the cost of information. Note that the problem can be rewritten as a choice of the signal probabilities π_s and posteriors q_s , instead of the signal structure p ; for any π_s and q_s such that $\sum_{s \in S} \pi_s q_s = q_0$, there is a unique signal structure p such that $\pi_s = \pi_s(p, q_0)$ and $q_s = q_s(p, q_0)$.

In the classic formulation of Sims, a problem of the form of (6) is considered, in which

¹⁴The full problem includes a choice over the signal alphabet S . A standard result, which will hold for all of the cost functions we study, is that $|S| = |A|$ is sufficient.

the cost function $C(p, q; S)$ is given by the mutual information between the signal and the state. Mutual information can be defined using Shannon's entropy,

$$H^{Shannon}(q) \equiv - \sum_{x \in X} (e_x^T q) \ln(e_x^T q), \quad (7)$$

where $e_x \in \mathbb{R}^X$ is the vector with a one corresponding to state x , and zeros elsewhere. Shannon's entropy can be used to define a measure of the degree to which each posterior q_s differs from the prior q_0 , the Kullback-Leibler (KL) divergence,

$$D_{KL}(q_s || q_0) \equiv H^{Shannon}(q_0) - H^{Shannon}(q_s) + (q_s - q_0)^T H_q^{Shannon}(q_0). \quad (8)$$

Mutual information is the expected value of the KL divergence over possible signals,

$$I^{Shannon}(p, q_0; S) \equiv \sum_{s \in S} \pi_s(p, q_0) D_{KL}(q_s(p, q_0) || q_0). \quad (9)$$

It is a measure of the informativeness of the signal structure p , in that it provides a measure of the degree to which the signal changes what one should believe about the state, on average. Mutual information is not, however, the only possible measure of the informativeness of an information structure, or the only plausible cost function for a static rational inattention problem. We introduce alternatives in the next section, but first return to our discussion of the continuous-time problem introduced in Section 2.

3.2 The Equivalence of Static and Dynamic Models

To prove our equivalence result, we restrict our attention to information-cost matrix functions that are "integrable," in the sense described by the following assumption.

Assumption 1. *There exists a twice-differentiable function $H : \mathbb{R}_+^{|X|} \rightarrow \mathbb{R}$ such that, for all*

q_t in the interior of the simplex,

$$Diag(q_t)^{-1}k(q_t)Diag(q_t)^{-1} = H_{qq}(q_t). \quad (10)$$

This class includes a number of information-cost matrix functions of interest: for example, it includes the case in which $k(q_t)$ is the inverse Fisher information matrix, which we will show corresponds to the standard rational inattention model, and the case in which $k(q_t)$ is the “neighborhood-based” function that we introduce in section 4. It is, however, more restrictive than the class of information-cost matrix functions defined in Section 2.¹⁵ We shall refer to the function H as the “entropy function,” for reasons that will become clear below. Note that $H(q)$ is convex, by the positive semi-definiteness of $k(q)$, and homogenous of degree one ($q^T \cdot H_{qq}(q) = \iota^T k(q)Diag(q)^{-1} = \vec{0}$).

For every convex function H , there is a “Bregman divergence,”

$$D_H(q_s||q) = H(q_s) - H(q) - (q_s - q)^T H_q(q). \quad (11)$$

The Kullback-Leibler divergence, for example, is a Bregman divergence (see (8)), with a entropy function equal to the negative of Shannon’s entropy.

To analyze our continuous time problem, we begin by proving that the information constraint (2) binds. Because the constraint binds, we can substitute the constraint into the HJB equation (3), and obtain the following result:

Lemma 1. *Anywhere the value function $V(q_t)$ is twice-differentiable and the DM chooses not to stop, for all $\sigma \in M(q_t)$,*

$$tr[\sigma^T \{Diag(q_t)V_{qq}(q_t)Diag(q_t) - \theta k(q_t)\}\sigma] \leq 0, \quad (12)$$

¹⁵It rules out, e.g., constant $k(q)$ (a hypothetical H would have asymmetric third-derivative cross-partials).

where $\theta = \chi^{-1} \kappa$, with equality under the optimal policy.

Proof. See the Appendix, Section B.1. □

The parameter θ , introduced in the lemma, describes the race between information acquisition and time in this model. The larger the penalty for delay, and the tighter the information constraint, the larger the parameter θ . We now describe our equivalence result, and then outline the key step of its proof, which relies on this lemma.

Theorem 1. *Under Assumption 1, the value function that solves the continuous time sequential evidence accumulation problem is*

$$V(q_0) = \max_{\pi \in \mathcal{P}(A), \{q_a \in \mathcal{P}(X)\}_{a \in A}} \sum_{a \in A} \pi(a) (u_a^T \cdot q_a) - \theta \sum_{a \in A} \pi(a) D_H(q_a || q_0),$$

subject to the constraint that $\sum_{a \in A} \pi(a) q_a = q_0$, where D_H is the Bregman divergence associated with the entropy function H that is defined by Assumption 1.

There exist maximizers π^* and q_a^* such that π^* is the unconditional probability, in the continuous time problem, of choosing a particular action, and q_a^* , for all a such that $\pi^*(a) > 0$, is the unique belief the DM will hold when stopping and choosing that action.

Proof. See the Appendix, Section section B.2. □

The continuous time sequential evidence accumulation problem is equivalent to a static rational inattention problem, with a particular kind of static information-cost function,

$$C(p, q_0; S) = \sum_{s \in S} \pi_s(p, q_0) D_H(q_s(p, q_0) || q_0), \quad (13)$$

with the signal space S identified with the set of possible actions A . Following Caplin et al. (2018b), we call such a cost function “uniformly posterior-separable.”

The mutual information cost function (9) proposed by Sims is one such cost function. In this case, the entropy function H is the negative of Shannon’s entropy (7), the corresponding information-cost matrix function is the inverse Fisher information matrix (4), the Bregman divergence is the Kullback-Leibler divergence (8), and the information measure defined by (13) is mutual information (9). Thus Theorem 1 provides a foundation for the standard static rational inattention model, and hence for the same predictions regarding stochastic choice as are obtained by Matějka et al. (2015).

On the other hand, Theorem 1 also implies that other cost functions can also be justified. Indeed, any (twice-differentiable) uniformly posterior-separable cost function (13) can be given such a justification, by choosing the information cost matrix function defined by equation (10). However, not all information cost matrix functions are reasonable. As discussed in the previous section, the information cost matrix function describes how hard it is to distinguish any pair of states. In many economic applications, there is a natural ordering or structure of the states, and we would like the information cost matrix function and the associated entropy function and Bregman divergence to reflect this structure. In the next section, we propose such a cost function.

First, however, we outline the key step of our proof, and then illustrate the important role that the information cost matrix function plays in comparative statics for the DM’s choices. Our proof strategy is best described as “guess and verify,” in that we start with the static value function described in Theorem 1 and then show that it is the value function of the continuous time model described in Section 2. The key step of the proof is to show that the static value function satisfies (12) in Lemma 1.¹⁶ For expositional purposes, we will

¹⁶Technical footnote: the “anywhere the value function is twice differentiable” caveat of Lemma 1 is relevant for our problem. The PDE described by equation (12) is “degenerate elliptic” and hence will not in general have a classical solution. Indeed, we do not prove that our static value function is twice differentiable everywhere, and suspect it is not at points where the “consideration set” (the set of actions with $\pi(a) > 0$, Caplin et al. (2018a)) changes. In our proof, we establish that the static problem value function is convex and continuously differentiable, which is sufficient to invoke a generalized version of Ito’s lemma for convex functions to verify that the static value function is the solution to the continuous time problem.

assume that the optimal policies of the static model, $\pi^*(a; q_0)$ and $q_a^*(q_0)$, are differentiable with respect to q_0 and strictly interior (we do not require these assumptions in the proof).

We begin by examining the first-order conditions with respect to q_a , and applying the envelope theorem. Let $\kappa(q_0)$ denote the vector of multipliers on the constraint that $\sum_{a \in A} \pi(a)q_a = q_0$. We have

$$\text{FOC: } u_a - \kappa(q_0) - \theta H_q(q_a^*(q_0)) + \theta H_q(q_0) = 0, \forall a \in A, \quad (14)$$

$$\text{ET: } V_q(q_0) = \kappa(q_0) + \sum_{a \in A} \pi^*(a; q_0)(q_a^*(q_0) - q_0)^T \cdot H_{qq}(q_0) = \kappa(q_0).$$

Now consider a perturbation $q_0 \rightarrow q_0 + \varepsilon z$, for some tangent vector z . Combining the FOC and ET, and then differentiating with respect to ε and evaluating at $\varepsilon = 0$,

$$V_{qq}(q_0) \cdot z = \theta H_{qq}(q_0) \cdot z - \theta H_{qq}(q_a^*(q_0)) \cdot \frac{dq_a^*(q_0 + \varepsilon z)}{d\varepsilon} \Big|_{\varepsilon=0}, \forall a \in A. \quad (15)$$

Observe that, due to the constraint, $\sum_{a \in A} \frac{d(\pi^*(a; q_0 + \varepsilon z)q_a^*(q_0 + \varepsilon z))}{d\varepsilon} \Big|_{\varepsilon=0} = z$. Multiplying both sides of equation (15) by $\frac{d(\pi^*(a; q_0 + \varepsilon z)q_a^*(q_0 + \varepsilon z)^T)}{d\varepsilon} \Big|_{\varepsilon=0}$, and then taking sums,

$$\begin{aligned} z^T \cdot V_{qq}(q_0) \cdot z &= \theta z^T \cdot H_{qq}(q_0) \cdot z \\ &\quad - \theta \sum_{a \in A} \pi^*(a; q_0) \left(\frac{dq_a^*(q_0 + \varepsilon z)}{d\varepsilon} \Big|_{\varepsilon=0} \right)^T \cdot H_{qq}(q_a^*(q_0)) \cdot \frac{dq_a^*(q_0 + \varepsilon z)}{d\varepsilon} \Big|_{\varepsilon=0} \\ &\quad - \theta \sum_{a \in A} \frac{d\pi^*(a; q_0 + \varepsilon z)}{d\varepsilon} \Big|_{\varepsilon=0} q_a^*(q_0)^T \cdot H_{qq}(q_a^*(q_0)) \cdot \frac{dq_a^*(q_0 + \varepsilon z)}{d\varepsilon} \Big|_{\varepsilon=0}. \end{aligned}$$

By Assumption 1, $q^T \cdot H_{qq}(q) = 0$, and hence the last line in this expression is zero. By the convexity of H , the summation on the second line is positive. Therefore, by Assumption 1,

$$z^T \cdot V_{qq}(q_0) \cdot z \leq z^T \cdot \text{Diag}(q_0)^{-1} k(q_0) \text{Diag}(q_0)^{-1} \cdot z,$$

establishing that (12) holds. To show that there is a direction z^* in which (12) holds with equality, it is sufficient to show that $\frac{dq_a^*(q_0 + \varepsilon z^*)}{d\varepsilon} \Big|_{\varepsilon=0} = 0$ for all $a \in A$. In any direction z spanned by the initial $q_a^*(q_0) - q_0$, it is not optimal for the DM to change the $q_a^*(q_0)$, only the probabilities $\pi^*(a, q_0)$ (this property, “Locally Invariant Posteriors,” was shown by Caplin et al. (2018b)). Thus, any of these directions can serve as z^* .¹⁷

3.3 Comparative Statics

Any uniformly posterior-separable cost function can be justified by some information cost matrix function. In this subsection, we ask whether the cost function matters in terms of the DM’s observable behavior. In some sense, the recoverability result of Caplin et al. (2018b) shows that the answer must be “yes”—if the cost function can be uniquely recovered from data on the likelihood of the DM’s action in each state, then that likelihood must be influenced by the cost function. Our particular point is to relate the comparative statics of the DM’s posteriors with respect to payoffs to the information cost matrix function.

To illustrate this point, we return to our running example of a buyer considering purchasing a security. The comparative statics of this case are particularly transparent, because there are only two actions, $A = \{L, R\}$; similar but more complicated formulas can be obtained in the many-action case. Combining the first-order conditions (14) for the two actions, and using the homogeneity of degree zero of $H_q(\cdot)$ (which follows from the homogeneity of degree one of H), we find that, continuing to assume interior solutions,

$$s - K\iota = \theta H_q(\pi_L^* q_L^*) - \theta H_q(q_0 - \pi_L^* q_L^*).$$

Now consider a perturbation of the security payoff for a particular asset value $x \in X$, $s(\varepsilon) =$

¹⁷That any such direction can serve as z^* indicates that there are (usually) many optimal policies in the continuous time problem that achieve the same value function. Intuitively, at each point, if the DM does not learn in some particular direction, she could always learn in that direction in the next instant.

$s + \varepsilon e_x$, and its effect on the probability of accepting the security conditional on some other $x' \in X$. Assuming that $q_L^* \neq q_0$, meaning that the buyer gathers some information,¹⁸ we can differentiate the above first-order equation and derive the comparative static

$$\frac{d(\pi_L^*(\varepsilon)e_{x'}^T q_L^*(\varepsilon))}{d\varepsilon} \Big|_{\varepsilon=0} = e_{x'}^T [\theta H_{qq}(\pi_L^* q_L^*) + \theta H_{qq}(q_0 - \pi_L^* q_L^*)]^{-1} e_x. \quad (16)$$

For this two-action model, this can be viewed as a definition of whether learning about the states x and x' are complements or substitutes (note that it is symmetric). If the reward for acceptance in state x increases, and learning about states x and x' are complements, the DM will endogenously choose to accept more often in state x' . By Assumption 1, the Hessians $H_{qq}(q)$ are transformed versions of the information cost matrix function $k(q)$. Intuitively, the difficulty of discriminating between the states x and x' , $-k_{xx'}(q)$, is closely related to whether or not they are complements or substitutes with respect to learning.

To understand the role of the matrix inverse, decompose the matrix $H_{qq}(\pi_L^* q_L^*) + H_{qq}(q_0 - \pi_L^* q_L^*)$ into a diagonal matrix $H_{diag}(\pi_L q_L)$ and the negative of an off-diagonal matrix $H_{off}(\pi_L q_L)$,

$$H_{qq}(\pi_L^* q_L^*) + H_{qq}(q_0 - \pi_L^* q_L^*) = H_{diag} - H_{off}.$$

Note that H_{diag} has entirely positive entries, and that the off-diagonal entries H_{off} are scaled versions of the off-diagonal entries of $k(q)$. We can write, ignoring the issue of whether the infinite sum converges,

$$[\theta H_{qq}(\pi_L^* q_L^*) + \theta H_{qq}(q_0 - \pi_L^* q_L^*)]^{-1} = \theta^{-1} H_{diag}^{-1} \sum_{j=0}^{\infty} (H_{diag}^{-1} H_{off})^j.$$

We can think of our comparative static (16) as consisting a sequence of “rounds” of effects. The “zero-round” effect is to increase the likelihood of the state x whose payoff has

¹⁸This condition ensures that $H_{qq}(\pi_L q_L) + H_{qq}(q_0 - \pi_L q_L)$ is positive-definite.

increased. The “first-round” effect is governed by the difficulty of discriminating between x and x' . If it is difficult to discriminate between x and x' , but easy to discriminate between x and x'' , the DM will want to make acceptance relatively more likely conditional on x' . However, this involves discriminating between x' and x'' , and hence leads to a “second-round” effect. Third- and higher-round effects can be defined in similar fashion. With the inverse Fisher information matrix as the cost function, these rounds of effects are a function only of $\pi_L q_L$ and q_0 and hence are identical for pairs (x, x') and (\tilde{x}, \tilde{x}') with identical probabilities under q_L and q_0 . With the example neighborhood-based information cost matrix function (5), first-round effects occur only between adjacent states, second round effects only between states that are “two spaces apart,” and so on.

At this point, we have shown that any uniformly posterior-separable cost function can be justified by our continuous time framework, with some information cost matrix function, that the information cost matrix function describes the difficulty or ease with which the DM can discern certain states, and that this difficulty or ease matters for behavior. Given these results, our next section proposes a specific class of information cost matrix functions, and hence uniformly posterior-separable cost functions, that we will argue are superior in certain respects to the standard mutual information cost function.

4 Neighborhood-Based Cost Functions

We begin by introducing a structure on the state space, to capture the idea that some pairs of states are harder to discern than others. Using this structure, we describe an information cost matrix function and the associated uniformly posterior-separable cost function.

4.1 The Neighborhood Framework

Suppose that the state space X can be written as the union of a finite collection of “neighborhoods” $\{X_i\}$, and let \mathcal{S} denote the set of these neighborhoods. Suppose furthermore that the state space is connected, in the sense that any two states can be connected by a sequence of overlapping neighborhoods.¹⁹ Define the selection matrices E_i as the $|X_i| \times |X|$ matrices that select each of the elements of X_i from a vector of length $|X|$.

For any belief $q \in \mathcal{P}(X)$, let $\bar{q}_i \equiv \sum_{x \in X_i} e_x^T q$ be the prior probability that some state belonging to neighborhood X_i occurs. Let $q_i \in \mathcal{P}(X_i)$ be the conditional probability distribution over states in neighborhood X_i , given the belief q and conditional on the state being in neighborhood X_i . That is, for all $x \in X_i$, $q_i \equiv \frac{1}{\bar{q}_i} E_i q$. The information cost matrix function has a neighborhood structure if, for all q in the interior of the simplex,

$$k_N(q; \rho) = \begin{cases} \sum_{i \in \mathcal{S}} c_i \bar{q}_i |X_i|^{1-\rho} E_i^T (g^+(q_i))^{2-\rho} E_i & \rho \neq 2, \\ \sum_{i \in \mathcal{S}} c_i \bar{q}_i |X_i|^{-1} E_i^T (I - q_i \mathbf{1}^T) (I - \mathbf{1} q_i^T) E_i & \rho = 2. \end{cases}$$

where $g^+(\cdot)$ is the inverse Fisher information matrix, $\rho \geq 1$ is a constant, and the constants c_i are strictly positive for each $i \in \mathcal{S}$. Within each neighborhood, we have assumed that the cost of information is proportional to the inverse Fisher information to some power, to retain to some of the tractability of the original rational inattention model. However, by varying the neighborhood structure, we can impose a variety of assumptions about how difficult it is to discriminate between various states. In the special case in which there is only a single neighborhood and $\rho = 1$, k_N is simply the inverse Fisher information matrix.

We should emphasize that using the inverse Fisher information matrix to some power within each neighborhood is useful from a tractability perspective, but is in no way essential

¹⁹That is, for any two states $x, x' \in X$, there exists a sequence of states $\{x_0, \dots, x_n\}$ with $x_0 = x, x_n = x'$, and the property that for any $1 \leq m \leq n$, states x_m and x_{m-1} belong to a common neighborhood.

to the idea of using neighborhoods. One could use a neighborhood structure with almost any “neighborhood” information cost matrix function. Our introduction of the exponent ρ (as opposed to simply using $\rho = 1$) follows Dean and Neligh (2018), who provide experimental evidence consistent with a neighborhood information structure, but find that values of ρ substantially larger than one provide the best fit to their experimental data. More generally, there is no reason aside from parsimony to assume that the same “neighborhood” information cost matrix function applies within each neighborhood.

We can use the information-cost matrix function in our continuous-time problem (the problem defined in Section 2).²⁰ It satisfies the Assumption 1, and hence Theorem 1 applies. That is, there is an entropy function $H_N(q)$ that can be used to define the static rational-inattention problem which is equivalent to the solution to the dynamic model.

Lemma 2. *Let $H^{Gen}(q_i; \rho)$ be the generalized entropy index of Shorrocks (1980) on the neighborhood $i \in \mathcal{I}$, defined for any interior q_i as*

$$H^{Gen}(q_i; \rho) = \begin{cases} \frac{1}{|X_i|} \frac{1}{(\rho-2)(\rho-1)} \sum_{x \in X_i} \{(|X_i| e_x^T q_i)^{2-\rho} - 1\} & \rho \notin \{1, 2\} \\ -\frac{1}{|X_i|} \sum_{x \in X_i} \ln(e_x^T q_i) & \rho = 2 \\ \sum_{x \in X_i} e_x^T q_i \ln(e_x^T q_i) & \rho = 1. \end{cases}$$

The entropy function $H_N(q; \rho)$ associated with the neighborhood-based information-cost matrix function $k_N(q; \rho)$ is, for any q in the relative interior of the simplex,

$$H_N(q; \rho) = - \sum_{i \in \mathcal{I}} c_i \bar{q}_i H^{Gen}(q_i; \rho),$$

and is defined on the boundary by continuity for $\rho < 2$ and as infinity for $\rho \geq 2$.

²⁰Our derivation of the continuous-time model (Hébert and Woodford (2018)) requires that $k_N(q) - mg^+(q) \geq 0$ for some strictly positive constant m . We can satisfy this requirement by including a neighborhood containing all states, with an arbitrarily small constant c_i associated with that neighborhood.

Proof. See the Appendix, Section B.3. □

The special case of $\rho = 1$ corresponds to (the negative of) Shannon's entropy within each neighborhood. The exponent ρ controls the curvature of the entropy function (Dean and Neligh (2018) use the following analogy: $H^{Gen}(q_i; \rho)$ is to Shannon's entropy as CRRA utility is to log utility). Using the our generalized entropy function $H_N(q; \rho)$, we can define a Bregman divergence, $D_N(q_s || q; \rho)$, as in (11), and a static rational inattention problem²¹ (as in Theorem 1),

$$V_N(q) = \max_{\pi \in \mathcal{P}(A), \{q_a \in \mathcal{P}(X)\}_{a \in A}} \sum_{a \in A} \pi(a) (u_a^T \cdot q_a) - \theta \sum_{a \in A} \pi(a) D_N(q_a || q; \rho). \quad (17)$$

It is sometimes more convenient to work with cost functions defined over signals $\{p_x \in \mathcal{P}(S)\}_{x \in X}$, as opposed to posteriors q_a and unconditional probabilities π (e.g. (6)). Below we rewrite the problem with a neighborhood cost function in this form, using Bayes' rule.

Lemma 3. *The static rational inattention problem in (17) can be written as*

$$\begin{aligned} V_N(q) = & \max_{\{p_x \in \mathcal{P}(S)\}_{x \in X}} \sum_{s \in S} \pi_s(p, q_0) \hat{u}(q_s(p, q_0)) \\ & - \theta \sum_{i \in \mathcal{I}} c_i |X_i|^{1-\rho} \bar{q}_i^{\rho-1} \sum_{x \in X_i} (e_x^T q)^{2-\rho} D_\rho(p_x || p E_i^T q_i), \end{aligned}$$

where

$$D_\rho(p_x || \pi) = \begin{cases} \frac{1}{(\rho-2)(\rho-1)} \sum_{s \in S: \pi_s > 0} \pi_s \left(\left(\frac{p_{x,s}}{\pi_s} \right)^{2-\rho} - 1 \right) & \rho \neq \{1, 2\} \\ \sum_{s \in S: \pi_s > 0} \pi_s \ln \left(\frac{\pi_s}{p_{x,s}} \right) & \rho = 2 \\ \sum_{s \in S: \pi_s > 0} p_{x,s} \ln \left(\frac{p_{x,s}}{\pi_s} \right) & \rho = 1. \end{cases}$$

Proof. See the Appendix, Section B.4. □

The divergences D_ρ are known as the α -divergences (under a different parameterization)

²¹To deal with the boundaries in the $\rho \geq 2$ case, we assume q has full support in this problem.

and are a transformed version of the Renyi divergences (Amari and Nagaoka (2007)). In the special case of $\rho = 1$, D_ρ is the Kullback-Leibler divergence. If $\rho = 1$ and there is only a single neighborhood, this is the standard rational inattention problem. The relevance of alternative neighborhood structures is illustrated by the following observation.

Corollary 1. *Consider a rational inattention problem with a neighborhood-based information-cost function (Lemma 3), and let x, x' be two states with the property that (i) $u_{a,x} = u_{a,x'}$ for all actions $a \in A$, (ii) $q_x = q_{x'}$, and (iii) the set of neighborhoods $\{X_i\}$ such that $x \in X_i$ is the same as the set such that $x' \in X_i$. Then under the optimal policy, $p_x^* = p_{x'}^*$. If $\rho = 1$, this result holds even if $q_x \neq q_{x'}$.*

Proof. The result follows directly from the problem in Lemma 3. □

The significance of Corollary 1 can be seen if we consider the predictions of rational inattention for a standard form of perceptual discrimination experiment, an application we describe in the next section. In these experiments, payments are based on correct and incorrect responses. As a result, two states in which the correct response and ex-ante likelihoods are identical will (for a single-neighborhood cost function) have the same likelihood of a correct response. Experimental evidence (intuitively) shows that in some states it is more difficult to determine the correct response than in other states.

4.2 A Specific Proposal: The Fisher Information Cost Function

Our neighborhood-based framework is flexible enough to accommodate a wide range of structures on the state space. However, in practice, we believe there is a particular structure that is relevant in many settings: when the states can be ordered on a line. Suppose there are $M + 1$ ordered states, $X^M = \{0, 1, \dots, M\}$, and that each pair of adjacent states forms a neighborhood, $X_i = \{i, i + 1\}$, for all $i \in \{0, 1, \dots, M - 1\}$. Thus two states belong to a common neighborhood if and only if one comes immediately after the other in the sequence.

This captures the idea that the available measurement technologies all respond similarly in states that are “similar,” in the sense of being at nearby positions in the sequence, so that repeated measurements are necessary to reliably distinguish between two states if and only if they are near each other in the sequence. Suppose further that $c_i = 1$ for all i , implying that it is equally difficult to distinguish two neighboring states at all points in the sequence.²²

Under these assumptions, for any q with full support,

$$H_N(q; \rho, M) = \frac{1}{\rho - 2} \frac{1}{\rho - 1} \sum_{j=0}^{M-1} \left(\frac{1}{2} (e_j^T + e_{j+1}^T) q \right) \times \left\{ \left(\frac{e_j^T q}{\frac{1}{2} (e_j^T + e_{j+1}^T) q} \right)^{2-\rho} + \left(\frac{e_{j+1}^T q}{\frac{1}{2} (e_j^T + e_{j+1}^T) q} \right)^{2-\rho} - 2 \right\}. \quad (18)$$

Defining the function $g(x; \rho) = \frac{1}{\rho-2} \frac{1}{\rho-1} x^{2-\rho}$, we can rewrite this expression as

$$H_N(q; \rho, M) = \sum_{j=0}^{M-1} \left(\frac{1}{2} (e_j^T + e_{j+1}^T) q \right) \left\{ g\left(1 - \frac{\frac{1}{2} (e_{j+1}^T - e_j^T) q}{\frac{1}{2} (e_j^T + e_{j+1}^T) q}; \rho\right) + g\left(1 + \frac{\frac{1}{2} (e_{j+1}^T - e_j^T) q}{\frac{1}{2} (e_j^T + e_{j+1}^T) q}; \rho\right) - 2g(1; \rho) \right\}.$$

This function penalizes differences in the function $g(\cdot; \rho)$ between states i and $i + 1$ and their average. Because the $g(\cdot; \rho)$ function is convex, any changes in probability are penalized. As a result, it will be optimal in the static rational inattention problem for the DM to smooth posterior probabilities across states of the world.

If q_i and q_{i+1} are close to each other for all i , a second-order Taylor approximation of the function $g(u; \rho)$ around $u = 1$ clarifies this point:

$$H_N(q; \rho, M) \approx \frac{1}{4} \sum_{j=0}^{M-1} \frac{((e_{j+1}^T - e_j^T) q)^2}{\frac{1}{2} (e_j^T + e_{j+1}^T) q}. \quad (19)$$

Note that this approximation is exact in the $\rho = 0$ case, and that the approximation is the

²²If c_i is the same for all i , we can without loss of generality set it equal to one, as the multiplier θ can still be used to scale the overall magnitude of information costs.

same for all values of ρ . Intuitively, all of the $H^{Gen}(q_i; \rho)$ resemble each other in the neighborhood of the uniform distribution, and hence when applied to a neighborhood with two states with similar probabilities are approximately identical.

For many applications, it is more convenient to work with a continuous state space. Based on this approximation result, it is tempting to infer that, in the limit as the number of states $M \rightarrow \infty$, if the discrete distributions q_M converge to differentiable function q ,

$$\lim_{M \rightarrow \infty} H_N(q_M; \rho, M) = \frac{1}{4} \int_{supp(q)} \frac{(q'(x))^2}{q(x)} dx,$$

where $supp(q)$ denotes the support of q . From this, we define a continuous state rational inattention problem:

$$\begin{aligned} V_N(q) = & \sup_{\pi \in \mathcal{P}(A), \{q_a \in \mathcal{P}_{LipG}\}_{a \in A}} \sum_{a \in A} \pi(a) \int_{supp(q)} u_a(x) q_a(x) dx \\ & - \frac{\theta}{4} \sum_{a \in A} \left\{ \pi(a) \int_{supp(q)} \frac{(q'_a(x))^2}{q_a(x)} dx \right\} + \frac{\theta}{4} \int_{supp(q)} \frac{(q'(x))^2}{q(x)} dx, \end{aligned} \quad (20)$$

subject to the constraint that, for all x ,

$$\sum_{a \in A} \pi(a) q_a(x) = q(x).$$

In this expression, x is the exogenous state, $u_a(x)$ is the utility of action $a \in A$ in state x , $q(x)$ is the prior over the states, and $q_a(x)$ is the posterior belief conditional on taking action a .

This problem has an alternative formulation as a choice of signal structure:

$$V_N(q) = \sup_{\{p_a\}_{a \in A} \in \mathcal{P}_{LipG}(A)} \int_{supp(q)} q(x) \sum_{a \in A} p_a(x) u_a(x) dx - \frac{\theta}{4} \int_{supp(q)} q(x) \sum_{a \in A} \frac{(p'_a(x))^2}{p_a(x)} dx, \quad (21)$$

where $\mathcal{P}_{LipG}(A)$ is the set of mappings $\{p_a : supp(q) \rightarrow \mathbb{R}_+\}_{a \in A}$ such that for each x ,

$\sum_{a \in A} p_a(x) = 1$, and for each action a , the function $p_a(x)$ is a differentiable function of x with a Lipschitz-continuous derivative.

This alternative formulation shows that our proposed static information-cost function is just the weighted (by $q(x)$) average of the Fisher information (Cover and Thomas (2012), sec. 11.10), a measure of the local discriminability of states.²³ It is for this reason that we refer to our proposal as the “Fisher information cost function.” Like mutual information, the Fisher-information cost function is a single-parameter cost function, and it can also be applied in almost any context, as long as the state space is continuous. Unlike mutual information, the Fisher information depends on the topological structure of the state space.

We prove the convergence of the static problem described in Theorem 1 to this problem formally in the Technical Appendix, Section C.2, under some regularity assumptions on the prior q (differentiability, with a Lipschitz-continuous derivative, and support on a compact set), for the specific case of $\rho = 1$.²⁴ In the proof, we show that the limiting optimal posteriors q_a are also differentiable and have the same support as q (so the Fisher information integrals make sense) and that their derivatives are also Lipschitz-continuous (which helps prove convergence). We refer to the set of full-support, differentiable probability distribution functions with Lipschitz-continuous derivatives as \mathcal{P}_{LipG} . The proof is quite technical, and the relevant economics are summarized by the approximation (19).

We have proposed a neighborhood structure that captures the idea that states might be ordered on a line. We now turn to applications, to illustrate the effects of using our proposed alternative in the place of the standard rational inattention cost function.

²³The equivalence of the two formulations is shown in the Technical Appendix, section C.2, where we also provide further discussion of the connection with Fisher information.

²⁴We also assume bounded utilities. We think the result holds for other values of ρ , and without some of our regularity assumptions, but generalizing our quite technical proof is difficult.

5 Applications of Neighborhood-Based Cost Functions

In this section, we apply our results to several applications of rational inattention. The first three applications concern binary choices, the fourth a linear-quadratic-Gaussian environment. These two environments cover a wide range of existing applications of rational inattention (for a survey, see Mackowiak et al. (2018)).

5.1 Security Design

In this application, we apply our framework to the mode of security design model with adverse selection in Yang (2017),²⁵ which builds on the buyer’s decision problem that we have used as an example thus far. We will briefly summarize the environment, and encourage readers to refer to Yang (2017) for a richer exposition.

A seller offers a security $s \in \mathbb{R}_+^{|X|}$, whose payoffs are contingent on the realized value of the assets backing the security, $x \in X \subseteq \mathbb{R}_+$, to a buyer at a price K . The buyer’s problem (our example earlier in the paper) is to gather information about which asset values $x \in X$ are most likely and then accept (“like,” L) or reject (R) this take-it-or-leave-it offer. Both parties are risk-neutral, and the seller discounts the cashflows by a factor $\beta < 1$, relative to the buyer. The security is constrained by limited liability, $0 \leq e_x^T s \leq x$.

The seller designs the security and offers a price, solving

$$\max_{s, K \geq 0} \pi_L(s, K) q_L(s, K)^T (K\mathbf{1} - \beta s).$$

subject to the limited liability constraint. In this expression, $\pi_L(s, K)$ and $q_L(x; s, K)$ are the optimal policies of the buyer who solves a rational inattention problem of the form of

²⁵Our neighborhood cost function could also be applied in the same fashion to the model of security design with moral hazard in attention described in the appendix of Hébert (2018).

Theorem 1, specialized to the case of two actions,

$$V(q_0; s, K) = \max_{\pi_L \in [0, 1], q_L, q_R \in \mathcal{P}(X)} \pi_L q_L^T (s - K\iota) - \theta \pi_L D_H(q_L || q_0) - \theta (1 - \pi_L) D_H(q_R || q_0),$$

subject to the constraint that $\pi_L q_L + (1 - \pi_L) q_R = q_0$.

Yang (2017) shows that, with the standard rational inattention cost function (D_H is the Kullback-Leibler divergence), the optimal security design is a debt contract, $s(x) = \min\{v(x), \bar{v}\}$ for some positive constant \bar{v} . The analysis involves two different cases, depending on whether the seller attempts to ensure acceptance with certainty ($\pi_L = 1$) or not, but the form of the optimal security is the same in both cases. To simplify our exposition, we will focus on the case with some possibility of rejection ($\pi_L < 1$), and discuss the case of acceptance with certainty in appendix section §C.3.

We explore, numerically, how the result of Yang (2017) changes with alternative $H(\cdot)$ functions. We consider three alternatives, our neighborhood-based function H_N with our pairwise neighborhood structure (equation (18)), a generalized entropy index cost function (the neighborhood cost function with only one neighborhood), and a “weighted” Shannon’s entropy. Weighted Shannon’s entropy (see, e.g., Nawrocki and Harding (1986)) is

$$H_w(q) = \sum_{x \in X} (e_x^T w) (e_x^T q) \ln\left(\frac{e_x^T q}{\iota^T q}\right),$$

where w is a vector of weights. Constant weights correspond to Shannon’s entropy.

Summarizing our results, we replicate numerically the proof of Yang (2017) that, with mutual information, the optimal security design is always a debt. In contrast, for weighted mutual information and the generalized entropy index, the shape of the security design depends on the weights and the prior, respectively. The neighborhood cost function, on the

other hand, appears to always generate the same shape irrespective of the prior, a result we speculate could be proven analytically in the continuous-state version of the model.

Below, we describe our calculation procedure, and the parameters we use to generate figures 3 and 4. In general, our choice of parameters is guided by a desire to illustrate the differences between the cost functions, and to ensure that acceptance is not certain ($\pi_L < 1$). Our numerical calculation uses the “first-order approach,” solving

$$\max_{s, K \geq 0, \pi_L \in [0, 1], q_L \in \mathcal{P}(X)} \pi_L q_L^T (K \mathbf{1} - \beta s)$$

subject to the buyer’s first order condition and that beliefs remain in the simplex,

$$\begin{aligned} s - K \mathbf{1} + \theta H_q(q_0 - \pi_L q_L) &= \theta H_q(\pi_L q_L), \\ e_x^T (q_0 - \pi_L q_L) &\geq 0, \forall x \in X, \end{aligned}$$

and the limited liability constraints.²⁶

Combining the first-order conditions of this security design problem and the constraints,

$$\begin{aligned} (1 - \beta)s^* &= \theta H_q(q - \pi_L^* q_L^*) - \theta H_q(\pi_L^* q_L^*) + \\ &+ \theta [H_{qq}(q - \pi_L^* q_L^*) + H_{qq}(\pi_L^* q_L^*)](\beta \pi_L^* q_L^* - \lambda + \nu), \end{aligned}$$

where λ and ν are the multipliers on the limited liability constraints. This illustrates that the optimal security design is determined by the entropy function, and hence the information cost matrix function, subject to the caveat that $\pi_L^* q_L^*$ is endogenous.²⁷

Our numerical experiment uses an X with twenty-one states, with values of x evenly spaced from 0 to 10. We use a seller β of 0.5, and prior q that is an equal-weighted mixture

²⁶We conjecture, but have not proven, that the first-order approach is valid in this context.

²⁷Our discussion of comparative statics in section §3 anticipates this result.

of a uniform and binomial (21 outcomes of a 50-50 coin flip) distribution. We have chosen these parameters to help illustrate the differences between the cost functions.²⁸

For the generalized entropy and neighborhood-based cost functions, we use $\rho = 13$. This value is close to the estimated parameter of Dean and Neligh (2018) for these two cost functions, although there is no particular reason to apply parameters estimated for perceptual experiments to security design. The various cost functions are not of the same “scale,” so the same values of θ do not necessarily result in the securities of the same scale. We have chosen $\theta = \frac{1}{2}$ for Shannon’s entropy, $\theta = 1$ for weighted Shannon’s entropy and the neighborhood cost function, and $\theta = \frac{1}{50}$ for the generalized entropy function, which results in securities that are of the same scale but distinct in our graphs.

For our weighted Shannon’s entropy, we use

$$w(x) = \frac{3}{2} + \frac{x}{10}.$$

This linear weight structure effectively assumes that it is more costly for the buyer to learn about good states than about bad states. We will see that this induces the seller to offer the buyer more in good states, and hence makes the buyer’s security more equity-like. The more general point, which we believe could be shown analytically, is that almost any security design could be reverse-engineered as optimal given some weight matrix. This reinforces the need to consider what kinds of information costs are reasonable.

Our numerical results are shown in figures 3 and 4. The first of these shows the optimal security designs, the second the optimal monotone (in x) security designs. Our numerical calculations recover the result of Yang (2017) for the case of Shannon’s entropy. They

²⁸In particular, the effects of weighted vs. standard Shannon’s entropy are proportional to $\ln(\beta)$, so we choose a value of β significantly different from one. The differences between the generalized entropy index and Shannon’s entropy disappear with a uniform prior, so we use the binomial part of the prior to highlight those differences. At the same time, it is helpful for numerical purposes to ensure the prior is significantly different from zero in each state, which is why we have the uniform part of the prior.

also illustrate our point that, with upward-sloping weights, the result for weighted Shannon’s entropy is equity-like. The “inverse hump-shape” of the optimal security with the generalized entropy index cost function is caused by the “hump-shape” of the prior.²⁹ The optimal securities for mutual information and weighted mutual information are monotone, and hence do not differ between the two graphs, whereas the optimal securities for the neighborhood based cost function and (imperceptibly) the generalized entropy index are non-monotone, and hence do differ. For weighted mutual information and the generalized entropy index, monotonicity or a lack thereof is not guaranteed, as the shape of the optimal security depends on the weights and prior, respectively.

Our results for the neighborhood cost function appear, regardless of parameters, to result in the same “debt-like,” but non-monotone, optimal security. This security is non-monotone and rapidly changing in one area. Rapid changes in security values would cause rapid changes in buyer behavior with Shannon’s entropy, and hence be sub-optimal, but this is not the case with neighborhood cost functions. As a result, it is possible for the optimal security to have rapid changes. However, when we restrict the security to be monotone, the optimal security is a debt, suggesting that the result of Yang (2017) is robust to using neighborhood cost functions (but not the other two alternatives) under this additional restriction. We conjecture that it is possible to prove the optimality of debt among monotone securities with a Fisher information cost, in the continuous state case.

Observant readers might notice a second feature of the optimal security for neighborhood-based cost functions: the “flat” part isn’t exactly flat. This feature arises from the “tri-diagonal” nature of the information cost matrix function $k(q)$, which leads to a difference equation describing the optimal security. As the number of states increases, the “flat” part of the security becomes increasingly flat. In the continuous state case, the difference equation becomes a differential equation, and we conjecture that the flat part is truly flat.

²⁹With a uniform prior, the optimal security with the generalized entropy index cost is also a debt.

5.2 Psychometric Functions

In this application, we discuss our theory in the context of perceptual experiments (for example, Shadlen et al. (2007) or Dean and Neligh (2018)). Suppose that the different states $X = \{0, 1, 2, \dots, M\}$, where M is an odd integer, represent different stimuli that may be presented to the subject, and that the subject is asked to classify the stimulus that is presented as one of two types (L or R); R is the correct answer if and only if $x > M/2$. For example, the stimuli might be visual images with different orientations relative to the vertical, with increasing values of x corresponding to increasingly clockwise orientations; the subject is asked whether the image is tilted clockwise or counter-clockwise relative to the vertical. The subject's goal is often simply to give as many correct responses as possible; hence we suppose that $u_{x,a} = 1$ if $a = R$ and $x > M/2$ or if $a = L$ and $x < M/2$, while $u_{x,a} = 0$ in all other cases. Each of the possible stimuli is presented with equal prior probability, and hence both responses have equal ex ante probability of being correct.

Both mutual information and generalizations based on the generalized entropy index correspond to a special case of a neighborhood-based cost function, in which all states belong to the unique neighborhood. Hence condition (iii) of Lemma 1 holds for any pair of states, and by assumption conditions (i) and (ii) hold as well. In the problem just posed, Lemma 1 implies that the probability of response R must be the same for all states $x < M/2$, and also the same (but higher) for all states $x > M/2$. Changing the severity of the information constraint changes the degree to which the probability of responding R is higher when $x > M/2$, but the response probabilities still will depend only on whether x is greater or less than $M/2$. This is illustrated in Appendix Figure 1, which plots the optimal response frequencies as a function of x , for alternative values of the cost parameter θ , in a numerical example in which C is given by mutual information and $M = 20$.

Alternatively, consider a neighborhood-based cost function in which the neighborhoods

are given by $X_i = \{i, i + 1\}$ for $i = 1, 2, \dots, M - 1$, and the constants c_i are equal to one for all neighborhoods, as in Section 4.2. Suppose further that $\rho = 1$ (noting that our approximation results suggest this choice does not matter very much). These assumptions suffice to completely determine a static information cost function (Lemma 2).

With this alternative neighborhood structure, Corollary 1 no longer requires that the response frequencies be identical for any two states. Moreover, because the cost function penalizes large differences in signal frequencies (and hence in response frequencies) in the case of neighboring states, in this case an optimal policy involves a gradual increase in the probability of response R as x increases, even though the payoffs associated with the different actions jump abruptly at a particular value of x . This is illustrated in Appendix Figure 2, which again shows the optimal response frequencies as a function of x , for alternative values of θ , in the case of the neighborhood cost function just described. The sigmoid functions predicted by rational inattention with this cost function — with the property that response frequencies differ only modestly from 50 percent when the stimuli are near the threshold of being correctly classified one way or the other, and yet approach zero or one in the case of stimuli that are sufficiently extreme — are characteristic of measured “psychometric functions” in perceptual experiments of this kind.³⁰

5.3 Global Games and The Fisher-Information Cost Function

The continuity of choice probabilities despite discrete changes in payoffs is also an important issue for the global games literature (Morris and Yang (2016)). This literature typically

³⁰For the general concept of a psychometric function, see, for example, Gabbiani and Cox (2010), chap. 25, especially Figures 25.1 and 25.2, and discussion on p. 360; or Gold and Heekeren (2014), p. 356. For an example of an empirical psychometric function for the kind of task discussed in the text (classification of the dominant direction of motion for a field of moving dots), see Shadlen et al. (2007), Figure 10.1A. Note not only that the curve is monotonically increasing, with many data points corresponding to different response probabilities between zero and one, but also that the subject’s reward function is clearly of the kind assumed in the text: only two possible reward levels (for correct vs. incorrect responses), with a discontinuous change in the reward where the sign of the “motion strength” changes from negative to positive.

assumes a continuum of states, so for this application we will discuss the continuous state limit described in (20). We will compare the Fisher information cost function we proposed in Section 4.2 with the more standard mutual information cost function.

This application is motivated by the work of Yang (2015) and Morris and Yang (2016), who study global games (e.g. Morris and Shin (1998)) with endogenous information acquisition. In the well-known analysis of Morris and Shin (1998), with exogenous private information, there is a unique equilibrium despite the incentives for coordination across DMs (subject to some caveats and details that are not relevant for our discussion). In contrast, Yang (2015) demonstrates that allowing for endogenous information acquisition, with mutual information as the information cost, restores a multiplicity of equilibria.

The key to Yang’s result is that DMs can tailor the signals they receive to sharply discriminate between nearby states of the world, as discussed in our previous example. As a result, they can all coordinate their decision (say, to invest or not) on a particular threshold, and there are many such thresholds that can represent equilibria if coordinated upon. But this result depends on the fact that the mutual-information cost function does not make it costly to have abrupt changes in signal probabilities as the state of the world changes continuously. Morris and Yang (2016) develop the complementary result, showing that even in the case of an endogenous information structure, if signal probabilities must vary continuously with the state, there is again a unique equilibrium.

Here we show that a neighborhood-based cost function can provide a justification for the kind of continuity condition that the result of Morris and Yang (2016) requires. Those authors study a global game with two possible actions, “invest” and “not-invest,” with equilibrium behavior characterized by a probability $s(x)$ of investing when the state is x . Their equilibrium uniqueness result depends on an assumption of continuous choice, meaning that for all information costs $\theta > 0$ and all parameterizations of the relevant utility function, $s(x)$ is absolutely continuous on a compact interval for which $q(x)$ has full support.

In our continuous state rational inattention problem (20), agents always choose posteriors that are differentiable, with a Lipschitz-continuous derivative. By assumption, the prior is also differentiable with a Lipschitz-continuous derivative. Therefore the function

$$s(x) = \frac{q_{invest}(x)}{q(x)} \pi_{invest}$$

is differentiable with respect to x in the support of q . By the Lipschitz-continuity of the derivatives $q'_{invest}(x)$ and $q'(x)$, and the fact that $q(x)$ has full support over the relevant compact interval, the derivative of $s(x)$ is bounded, and hence $s(x)$ is absolutely continuous.

Thus, our proposal provides a micro-foundation for the continuous choice assumption required by Morris and Yang (2016), and hence for uniqueness in global games.

5.4 Linear-Quadratic Gaussian Environments

In this application, we consider the classic “Linear-Quadratic-Gaussian” (LQG) tracking problem, which is a major application of the standard theory of rational inattention (see, e.g., Sims (2010)). To consider this application, we extend the continuous-state version of our model, with the Fisher information cost function, to a continuous action space (we do not formally prove convergence). The message of this application, unlike the message of our three previous applications, is that the model predicts the same behavior regardless of whether the information cost is mutual information or Fisher information.

Let the state space X be the real line, and let the space of possible actions A be the real line as well. We assume that the DM’s task is to estimate the value of the state (i.e., to “track” variation in the state), with a reward given by $u_a(x) = -(x - a)^2$. In other words, the goal is to minimize the mean squared error of the DM’s estimate.

We assume that the prior distribution over the state space X is a Gaussian distribution, with mean μ and variance σ^2 . We assume that information costs are given by the Fisher-

information cost function, which we now generalize to allow for a continuum of actions.³¹

It is convenient to consider the conditional probabilities $\{p_a(x)\}_{a \in A}$, as in (21). Our problem is to choose the functions $\{p_a(x)\}_{a \in A}$ so as to minimize

$$\int_{-\infty}^{\infty} q(x) \int_{-\infty}^{\infty} [p_a(x)(a-x)^2 + \frac{\theta}{4} \frac{(p'_a(x))^2}{p_a(x)}] da dx,$$

subject to the constraint that $\sum_{a \in A} p_a(x) = 1$ for all $x \in X$.

This is a problem in the calculus of variations. We show in Technical Appendix, Section C.4, that if $\theta < 4\sigma^2$, the optimal information structure is equivalent to observing a noisy signal $s = x + \varepsilon$, with the “measurement error” $\varepsilon \sim N(0, v^2)$. Consequently, defining $\beta = \frac{\sigma^2}{\sigma^2 + v^2}$, the DM’s posterior mean (and optimal action) is

$$E[x|s] = (1 - \beta)\mu + \beta s,$$

Moreover, the optimal degree of noise in the signal s is given by $v^2 = \sigma^2[2\sigma^2\theta^{-\frac{1}{2}} - 1]^{-1}$, which is an increasing function of θ for all $\theta < 4\sigma^2$.

In this solution, $p_a(x)$ is a normal density function for each value of x , with a mean that is a linear function of x , and a variance that is independent of x . As θ approaches the bound $4\sigma^2$, the optimal value of v^2 grows without bound, and β approaches zero; in the limit, the information structure becomes perfectly uninformative. We further show in the technical appendix that for any $\theta \geq 4\sigma^2$, it is optimal for the information structure to be purely uninformative, and for the DM to choose an action $a = \mu$ regardless of the state. Therefore, this model, like the rational inattention model of Sims, allows for the possibility of a corner solution in which there is no attention at all paid to some features of the environment, despite the fact that tracking them would allow the DM to achieve a

³¹We also generalize the model use the entire real line instead of a bounded interval as the state space.

higher level of welfare, and despite a finite information cost parameter θ .

More generally, the main features of our results are exactly the same those of the LQG tracking problem with a mutual information. We include these results to show that, if one considers the tractability of the LQG problem an appealing feature of mutual information, the problem remains equally tractable (and the results equally sensible) with the Fisher information cost function.

6 Conclusion

What kinds of information cost functions should be used in static rational inattention problems? We have argued that two particular, related properties are desirable. First, the cost function should represent the results of a sequential evidence accumulation problem, one that can be related to the existing literature in psychology and neuroscience. Second, the cost function should capture the idea that certain states are easier or harder to discriminate than others. These two properties are linked, in our continuous time model, by the information cost matrix function, which controls both the ultimate choice probabilities (via the entropy function) and describes the difficulty of discriminating between pairs of states.

We have shown that all uniformly posterior separable cost functions satisfy the first property, and we have introduced the neighborhood-based cost functions as a subset of that also satisfy the second property. We have also extended our model to the continuous state case, and shown that Fisher information is the continuous state analog of the cost functions we advocate. In most of our applications, our proposed cost functions and mutual information generate very different predictions, but (reassuringly) they generate the same predictions for the linear-quadratic-Gaussian problem. We believe these neighborhood-based cost functions can and should be used in almost all applications of the rational inattention theory in the place of mutual information.

References

- Shun-ichi Amari and Hiroshi Nagaoka. *Methods of information geometry*, volume 191. American Mathematical Soc., 2007.
- Andrew Caplin and Mark Dean. The behavioral implications of rational inattention with Shannon entropy. *Unpublished manuscript*, August 2013.
- Andrew Caplin, Mark Dean, and John Leahy. Rational inattention, optimal consideration sets and stochastic choice. Technical report, 2018a.
- Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *Unpublished manuscript*, 2018b.
- Yeon-Koo Che and Konrad Mierendorff. Optimal sequential decision with limited attention. *Unpublished manuscript*, 2017.
- John A Clithero. Improving out-of-sample predictions using response times and a model of the decision process. *Journal of Economic Behavior & Organization*, 148:344–375, 2018.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Henrique De Oliveira, Tommaso Denti, Maximilian Mihm, and Kemal Ozbek. Rationally inattentive preferences and hidden information costs. *Theoretical Economics*, 12:621–624, 2017.
- Mark Dean and Nathaniel Neligh. Experimental tests of rational inattention. Technical report, Working Paper, Columbia University, 2018.
- Ambuj Dewan and Nate Neligh. Estimating information cost functions in models of rational inattention. *Unpublished manuscript*, January 2017.
- Ernst Fehr and Antonio Rangel. Neuroeconomic foundations of economic choice — recent advances. *Journal of Economic Perspectives*, 25(4):3–30, 2011.
- Drew Fudenberg, Philipp Strack, and Tomasz Strzalecki. Speed, accuracy, and the optimal timing of choices. *American Economic Review*, forthcoming.
- Fabrizio Gabbiani and Steven J. Cox. *Mathematics for Neuroscientists*. Academic Press, 2010.
- Joshua I. Gold and Hauke R. Heekeren. Neural mechanisms for perceptual decision making. In Paul W. Glimcher and Ernst Fehr, editors, *Neuroeconomics: Decision Making and the Brain*, 2d ed. Academic Press, 2014.
- Benjamin Hébert. Moral hazard and the optimality of debt. *The Review of Economic Studies*, 85(4):2214–2252, 2018.
- Benjamin Hébert and Michael Woodford. Rational inattention in continuous time. *Unpublished manuscript*, September 2018.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Ian Krajbich, Bastiaan Oud, and Ernst Fehr. Benefits of neuroeconomics modeling: New policy interventions and predictors of preference. *American Economic Review*, 104(5):

- 501–506, 2014.
- Bartosz Mackowiak, Filip Matejka, and Mirko Wiederholt. Rational inattention: A disciplined behavioral model. 2018.
- Filip Matějka, Alisdair McKay, et al. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.
- Stephen Morris and Hyun Song Shin. Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, pages 587–597, 1998.
- Stephen Morris and Philipp Strack. The Wald problem and the equivalence of sequential sampling and static information costs. *Unpublished manuscript*, June 2017.
- Stephen Morris and Ming Yang. Coordination and the relative cost of distinguishing nearby states. *Unpublished manuscript*, 2016.
- Giuseppe Moscarini and Lones Smith. The optimal level of experimentation. *Econometrica*, 69(6):1629–1644, 2001.
- David N Nawrocki and William H Harding. State-value weighted entropy as a measure of investment risk. *Applied Economics*, 18(4):411–419, 1986.
- R Tyrrell Rockafellar. *Convex analysis*. 1970.
- Michael Shadlen and Daphna Shohamy. Decision making and sequential sampling from memory. *Neuron*, 90(5):927–939, 2016.
- Michael N. Shadlen et al. The speed and accuracy of a perceptual decision: A mathematical primer. In K. Doya et al., editors, *Bayesian Brain: Probabilistic Approaches to Neural Coding*. M.I.T. Press, 2007.
- Anthony F Shorrocks. The class of additively decomposable inequality measures. *Econometrica: Journal of the Econometric Society*, pages 613–625, 1980.
- Christopher A Sims. Rational inattention and monetary economics. *Handbook of Monetary Economics*, 3:155–181, 2010.
- Satohiro Tajima, Jan Drugowitsch, and Alexandre Pouget. Optimal policy for value-based decision-making. *Nature communications*, 7, 2016.
- Cédric Villani. *Topics in optimal transportation*. Number 58. American Mathematical Soc., 2003.
- Michael Woodford. Inattentive valuation and reference-dependent choice. *Unpublished manuscript*, May 2012.
- Michael Woodford. An optimizing neuroeconomic model of discrete choice. Technical report, National Bureau of Economic Research, February 2014.
- Ming Yang. Coordination with flexible information acquisition. *Journal of Economic Theory*, 158:721–738, 2015.
- Ming Yang. Optimality of debt under flexible information acquisition. 2017.
- Weijie Zhong. Optimal dynamic information acquisition. *Unpublished manuscript*, September 2018.

A Figures




Figure 1: Predicted response probabilities with a mutual-information cost function, for alternative values of the cost parameter θ .




Figure 2: Predicted response probabilities with a neighborhood-based cost function, in which each neighborhood consists only of two adjacent states.




Figure 3: Optimal Security Designs by Entropy Function




Figure 4: Optimal Monotone Security Designs by Entropy Function

B Proofs

B.1 Proof of Lemma 1

In the continuation region, everywhere the value function is twice differentiable,

$$\sup_{\sigma_t \in M(q_t)} \frac{1}{2} \text{tr}[\sigma_t^T D(q_t) V_{qq}(q_t) D(q_t) \sigma_t] = \kappa,$$

subject to

$$\frac{1}{2} \text{tr}[\sigma_t^T k(q_t) \sigma_t] \leq \chi.$$

First, suppose that the constraint does not bind and a maximizing optimal policy exists:

$$\frac{1}{2} \text{tr}[\sigma_t^{*T} k(q_t) \sigma_t^*] = a\chi,$$

where σ_t^* is a maximizer, for some $a \in [0, 1)$ ($a \geq 0$ by the positive semi-definiteness of $k(q_t)$). For any $c \in (1, a^{-1})$, with $a^{-1} = \infty$ for $a = 0$, if we used $\sigma_t = c\sigma_t^*$ instead, the policy would be feasible and we would have

$$\frac{1}{2} \text{tr}[\sigma_t^T D(q_t) V_{qq}(q_t) D(q_t) \sigma_t] = c^2 \kappa > \frac{1}{2} \text{tr}[\sigma_t^{*T} D(q_t) V_{qq}(q_t) D(q_t) \sigma_t^*] = \kappa,$$

a contradiction by the fact that $\kappa > 0$. Therefore, either the constraint binds under the optimal policy or an optimal policy does not exist. The latter would require that, for some non-zero vector $z \in \mathbb{R}^{|X|}$ with $zz^T \in M(q_t)$,

$$z^T D(q_t) V_{qq}(q_t) D(q_t) z > 0$$

and $z^T k(q_t) z = 0$, but the null space of $k(q_t)$ consists only of vectors whose elements are constant over the support of q_t , and therefore satisfy $q^T z \neq 0$, implying that $zz^T \notin M(q_t)$. Therefore, the constraint binds, and an optimal policy exists.

Using θ as defined in the lemma, it must be the case (anywhere the DM chooses not to stop and the value function is twice differentiable) that

$$\max_{\sigma_t \in M(q_t)} \frac{1}{2} \text{tr}[\sigma_t \sigma_t^T (\text{Diag}(q_t) V_{qq}(q_t) \text{Diag}(q_t) - \theta k(q_t))] = 0.$$

B.2 Proof of Theorem 1

Define $\phi(q_t)$ as the static value function in the statement of the theorem (we will prove that it is equal to $V(q_t)$, the value function of the dynamic problem). We first show that $\phi(q_t)$ satisfies the HJB equation, can be implemented by a particular strategy for the DM, and that any other strategy for the DM achieves weakly less utility. We begin by observing that

$$\iota^T k(q_t) \text{Diag}(q_t)^{-1} = 0 = \iota^T \text{Diag}(q_t) H_{qq}(q_t) = q_t^T H_{qq}(q_t),$$

and therefore converse of Euler's homogenous function theorem applies. That is, $H_q(q_t)$ is homogenous of degree zero, and $H(q_t)$ is homogeneous of degree one.

We start by showing that the function $\phi(q_t)$ is twice-differentiable in certain directions. Substituting the definition of the divergence into the statement of theorem,

$$\phi(q_0) = \max_{\pi \in \mathcal{P}(A), \{q_a \in \mathcal{P}(X)\}_{a \in A}} \sum_{a \in A} \pi(a) u_a^T \cdot q_a + \theta H(q_0) - \theta \sum_{a \in A} \pi(a) H(q_a),$$

subject to the same constraint. Define a new choice variable, $\hat{q}_a = \pi(a) q_a$. By definition, $\hat{q}_a \in \mathbb{R}_+^{|X|}$, and the constraint is $\sum_{a \in A} \hat{q}_a = q_0$. By the homogeneity of H , the objective is

$$\phi(q_0) = \max_{\pi \in \mathcal{P}(A), \{q_a \in \mathcal{P}(X)\}_{a \in A}, \{\hat{q}_a \in \mathcal{P}(X)\}_{a \in A}} \sum_{a \in A} u_a^T \cdot \hat{q}_a + \theta H(q_0) - \theta \sum_{a \in A} H(\hat{q}_a).$$

Any choice of \hat{q}_a satisfying the constraint can be implemented by some choice of π and q_a in the following way: set $\pi(a) = \iota^T \hat{q}_a$, and (if $\pi(a) > 0$) set

$$q_a = \frac{\hat{q}_a}{\pi(a)}.$$

If $\pi(a) = 0$, set $q_a = q_0$. By construction, the constraint will require that $\pi(a) \leq 1$, $\sum_{a \in A} \pi(a) = 1$, and the fact that the elements of q_a are weakly positive will ensure $\pi(a) \geq 0$. Similarly, $\iota^T q_a = 1$ for all $a \in A$, and the elements of q_a are weakly greater than zero. Therefore, we can implement any set of \hat{q}_a satisfying the constraints.

Rewriting the problem in Lagrangian form,

$$\begin{aligned} \phi(q_0) = & \max_{\{\hat{q}_a \in \mathbb{R}^{|X|}\}_{a \in A}} \min_{\kappa \in \mathbb{R}^{|X|}, \{v_a \in \mathbb{R}_+^{|X|}\}_{a \in A}} \sum_{a \in A} u_a^T \cdot \hat{q}_a + \theta H(q_0) \\ & - \theta \sum_{a \in A} H(\hat{q}_a) + \kappa^T (q_0 - \sum_{a \in A} \hat{q}_a) + \sum_{a \in A} v_a^T \hat{q}_a. \end{aligned}$$

Observe that $\phi(q_0)$ is convex in q_0 . Suppose not: for some $q = \lambda q_0 + (1 - \lambda)q_1$, with $\lambda \in (0, 1)$, $\phi(q) < \lambda\phi(q_0) + (1 - \lambda)\phi(q_1)$. Consider a relaxed version of the problem in which the DM is allowed to choose two different \hat{q}_a for each a . Because of the convexity of H , even with this option, the DM will set both of the \hat{q}_a to the same value, and therefore the relaxed problem reaches the same value as the original problem. However, in the relaxed problem, choosing the optimal policies for q_0 and q_1 in the original problem, scaled by λ and $(1 - \lambda)$ respectively, is feasible. It follows that $\phi(q) \geq \lambda\phi(q_0) + (1 - \lambda)\phi(q_1)$. Note also that $\phi(q_0)$ is bounded on the interior of the simplex. It follows by Alexandrov's theorem that ϕ is twice-differentiable almost everywhere on the interior of the simplex.

By the convexity of H , the objective function is concave, and the constraints are affine and a feasible point exists. Therefore, the KKT conditions are necessary. Anywhere the objective function is continuously differentiable in the choice variables and in q_0 , and therefore the envelope theorem applies. We have, by the envelope theorem,

$$\phi_q(q_0) = \theta H_q(q_0) + \kappa,$$

and the first-order conditions (for all $a \in A$ with $\hat{q}_a \neq \vec{0}$),

$$u_a - \theta H_q(\hat{q}_a) - \kappa + v_a = 0. \quad (22)$$

If $\hat{q}_a = \vec{0}$, we must have $q^T(u_a - \kappa) \leq \theta H(q)$ for all q , meaning that $u_a - \kappa$ is a sub-gradient of $H(q)$ at $q = 0$. In this case, we can define $v_a = \vec{0}$ and observe that the first-order condition holds for an appropriately-chosen sub-gradient. Define $\hat{q}_a(q_0)$, $\kappa(q_0)$, and $v_a(q_0)$ as functions that are solutions to the first-order conditions and constraints.

We next prove the ‘‘locally invariant posteriors’’ property described by Caplin et al. (2018b). Consider an alternative prior, $\tilde{q}_0 \in \mathcal{P}(X)$, such that

$$\tilde{q}_0 = \sum_{a \in A} \alpha(a) \hat{q}_a(q_0)$$

for some $\alpha(a) \geq 0$. Conjecture that $\hat{q}_a(\tilde{q}_0) = \alpha(a) \hat{q}_a(q_0)$, $\kappa(\tilde{q}_0) = \kappa(q_0)$, and $v_a(\tilde{q}_0) = v_a(q_0)$. By the homogeneity property,

$$H_q(\alpha(a) \hat{q}_a(q_0)) = H_q(\hat{q}_a(q_0)),$$

and therefore the first-order conditions are satisfied. By construction, the constraint is

satisfied, the complementary slackness conditions are satisfied, and \hat{q}_a and v_a are weakly positive. Therefore, all necessary conditions are satisfied, and by the concavity of the problem, this is sufficient. It follows that the conjecture is verified.

Consider a perturbation

$$q_0(\varepsilon; z) = q_0 + \varepsilon z,$$

with $z \in \mathbb{R}^{|X|}$, such that $q_0(\varepsilon; z)$ remains in $\mathcal{P}(X)$ for some $\varepsilon > 0$. If z is in the span of $\hat{q}_a(q_0)$, then there exists a sufficiently small $\varepsilon > 0$ such that the above conjecture applies. In this case that κ is constant, and therefore $\phi_q(q_0(\varepsilon; z))$ is directionally differentiable with respect to ε . If $q_0(-\varepsilon; z) \in \mathcal{P}(X)$ for some $\varepsilon > 0$, then ϕ_q is differentiable, with

$$\phi_{qq}(q_0) \cdot z = \theta H_{qq}(q_0) \cdot z,$$

proving twice-differentiability in this direction. This perturbation exists anywhere the span of $\hat{q}_a(q_0)$ is strictly larger than the line segment connecting zero and q_0 (in other words, all $\hat{q}_a(q_0)$ are not proportional to q_0). Define this region as the continuation region, Ω . Outside of this region, all $\hat{q}_a(q_0)$ are proportional to q_0 , implying that

$$\phi(q_0) = \max_{a \in A} u_a^T \cdot q_0,$$

as required for the stopping region. Within the continuation region, the strict convexity of $H(q_0)$ in all directions orthogonal to q_0 implies that, as required,

$$\phi(q_0) > \max_{a \in A} u_a^T \cdot q_0.$$

Now consider an arbitrary perturbation z such that $q_0(\varepsilon; z) \in \mathbb{R}_+^{|X|}$ and $q_0(-\varepsilon; z) \in \mathbb{R}_+^{|X|}$ for some $\varepsilon > 0$. Observe that, by the constraint,

$$\varepsilon z = \sum_{a \in A} (\hat{q}_a(\varepsilon; z) - \hat{q}_a(q_0)).$$

It follows that

$$(\kappa^T(q_0(\varepsilon; z)) - \kappa^T(q_0))\varepsilon z = \sum_{a \in A} (\kappa^T(q_0(\varepsilon; z)) - \kappa^T(q_0))(\hat{q}_a(\varepsilon; z) - \hat{q}_a(q_0)).$$

By the first-order condition,

$$(\kappa^T(q_0(\varepsilon; z)) - \kappa^T(q_0))(\hat{q}_a(\varepsilon; z) - \hat{q}_a(q_0)) = [\theta H_q(\hat{q}_a(q_0)) - \theta H_q(\hat{q}_a(\varepsilon; z)) + \mathbf{v}_a^T(q_0(\varepsilon; z)) - \mathbf{v}_a^T(q_0)](\hat{q}_a(\varepsilon; z) - \hat{q}_a(q_0)).$$

Consider the term

$$(\mathbf{v}_a^T(q_0(\varepsilon; z)) - \mathbf{v}_a^T(q_0))(\hat{q}_a(\varepsilon; z) - \hat{q}_a(q_0)) = \sum_{x \in X} (\mathbf{v}_a^T(q_0(\varepsilon; z)) - \mathbf{v}_a^T(q_0)) e_x e_x^T (\hat{q}_a(\varepsilon; z) - \hat{q}_a(q_0)).$$

By the complementary slackness condition,

$$(\mathbf{v}_a^T(q_0(\varepsilon; z)) - \mathbf{v}_a^T(q_0))(\hat{q}_a(\varepsilon; z) - \hat{q}_a(q_0)) = -\mathbf{v}_a^T(q_0(\varepsilon; z))\hat{q}_a(q_0) - \mathbf{v}_a^T(q_0)\hat{q}_a(\varepsilon; z) \leq 0.$$

By the convexity of H ,

$$\theta(H_q(\hat{q}_a(q_0)) - \theta H_q(\hat{q}_a(\varepsilon; z)))(\hat{q}_a(\varepsilon; z) - \hat{q}_a(q_0)) \leq 0.$$

Therefore,

$$(\kappa^T(q_0(\varepsilon; z)) - \kappa^T(q_0))\varepsilon z \leq 0.$$

Thus, anywhere ϕ is twice differentiable (almost everywhere on the interior of the simplex),

$$\phi_{qq}(q) \preceq \theta H_{qq}(q),$$

with equality in certain directions. Therefore, it satisfies the HJB equation almost everywhere in the continuation region. Moreover, by the convexity of ϕ ,

$$(H_q(q_0(\varepsilon; z)) - H_q(q_0))^T \varepsilon z \geq (\phi_q(q_0(\varepsilon; z)) - \phi_q(q_0))^T \varepsilon z \geq 0,$$

implying that the ‘‘Hessian measure’’ (see Villani (2003)) associated with ϕ_{qq} has no pure point component. This implies that ϕ is continuously differentiable.

Next, we show that there is a strategy for the DM in the dynamic problem which can implement this value function. Suppose the DM starts with beliefs q_0 , and generates some $\hat{q}_a(q_0)$ as described above. As shown previously, this can be mapped into a policy $\pi(a, q_0)$

and $q_a(q_0)$, with the property that

$$\sum_{a \in A} \pi(a, q_0) q_a(q_0) = q_0.$$

We will construct a policy such that, for all times t ,

$$q_t = \sum_{a \in A} \pi_t(a) q_a(q_0)$$

for some $\pi_t(a) \in \mathcal{P}(A)$. Let Ω (the continuation region) be the set of q_t such that a $\pi_t \in \mathcal{P}(A)$ satisfying the above property exists and $\pi_t(a) < 1$ for all $a \in A$. The associated stopping rule will be the stop whenever $\pi_t(a) = 1$ for some $a \in A$.

For all $q_t \in \Omega$, there is a linear map from $\mathcal{P}(A)$ to Ω , which we will denote $Q(q_0)$:

$$Q(q_0)\pi_t = q_t.$$

Therefore, we must have

$$Q(q_0)d\pi_t = \text{Diag}(q_t)\sigma_t dB_t.$$

By the assumption that $|X| \geq |A|$, there exists a $|A| \times |X|$ matrix $\sigma_{\pi,t}$ such that

$$Q(q_0)\sigma_{\pi,t} = \text{Diag}(q_t)\sigma_t$$

and $d\pi_t = \sigma_{\pi,t} dB_t$. Define $\tilde{\phi}(\pi_t) = \phi(q_t)$. As shown above,

$$Q^T(q_0)\phi_{qq}(q_t)Q(q_0)$$

exists everywhere in Ω , and therefore

$$\tilde{\phi}(\pi_t) - \theta H(Q(q_0)\pi_t)$$

is a martingale. We also have to scale $\sigma_{\pi,t}$ to respect the constraint,

$$\frac{1}{2} \text{tr}[\sigma_t \sigma_t^T k(q_t)] = \chi > 0.$$

This can be rewritten as

$$\frac{1}{2}tr[\sigma_{\pi,t}\sigma_{\pi,t}^T Q^T(q_0)Diag^+(Q(q_0)\pi_t)k(Q(q_0)\pi_t))Diag^+(Q(q_0)\pi_t)Q(q_0)] = \chi,$$

where $Diag^+$ denotes the pseudo-inverse of the diagonal matrix.

By the positive-definiteness of k in all directions except those constant in the support of $Q(q_0)\pi_t$, we will always have $\frac{1}{2}tr[\sigma_{\pi,t}\sigma_{\pi,t}^T] > 0$. Under the stopping rule described previously, the boundary will be hit a.s. as the horizon goes to infinity. As a result, by the martingale property described above, initializing $\pi_0(a) = \pi(a, q_0)$,

$$\tilde{\phi}(\pi_0) = E_0[\tilde{\phi}(\pi_\tau) - \theta H(Q(q_0)\pi_\tau) + \theta H(Q(q_0)\pi_0)].$$

By Ito's lemma,

$$\theta H(Q(q_0)\pi_\tau) - \theta H(Q(q_0)\pi_0) = \int_0^\tau \chi \theta dt = \mu \tau.$$

By the value-matching property of ϕ , $\tilde{\phi}(\pi_\tau) = \hat{u}(Q(q_0)\pi_\tau)$. It follows that, as required,

$$\phi(q_0) = \tilde{\phi}(\pi_0) = E_0[\hat{u}(q_\tau) - \mu \tau].$$

Finally, we verify that alternative policies are sub-optimal. Consider an arbitrary control process σ_t and stopping rule described by the stopping time τ . We have, by the convexity of ϕ and the generalized Ito formula for convex functions (noting that we have shown that the Hessian measure associated with ϕ_{qq} has no pure point component), interpreting ϕ_{qq} in a distributional sense,

$$E_0[\phi(q_\tau)] - \phi(q_0) = \frac{1}{2}E_0\left[\int_0^\tau tr[\sigma_t^T D(q_t)\phi_{qq}(q_t)D(q_t)\sigma_t]dt\right].$$

By the feasibility of the policies, anywhere in the continuation region of the optimal policy,

$$\frac{1}{2}tr[\sigma_t^T D(q_t)\phi_{qq}(q_t)D(q_t)\sigma_t] \leq \frac{1}{2}\theta tr[\sigma_t^T k(q_t)\sigma_t] \leq \theta \chi.$$

In the stopping region of the optimal policy,

$$\frac{1}{2}tr[\sigma_t^T D(q_t)\phi_{qq}(q_t)D(q_t)\sigma_t] = 0 < \theta \chi.$$

Therefore,

$$\phi(q_0) \geq E_0[\phi(q_\tau)] - \int_0^\tau \theta \chi dt.$$

By inequality $\phi(q_\tau) \geq \hat{u}(q_\tau)$, $\phi(q_0) \geq E_0[\hat{u}(q_\tau) - \mu \tau]$ for all policies, verifying optimality.

B.3 Proof of Lemma 2

We have, for any interior q ,

$$\begin{aligned} H_N(q; \rho) &= - \sum_{i \in \mathcal{I}} c_i \bar{q}_i H^{Gen}(q_i; \rho) \\ &= \sum_{i \in \mathcal{I}} c_i \bar{q}_i \frac{1}{|X_i|} \frac{1}{(\rho-2)(\rho-1)} \sum_{x \in X_i} \left\{ \left(\frac{e_x^T q}{\frac{1}{|X_i|} \bar{q}_i} \right)^{2-\rho} - 1 \right\}. \end{aligned}$$

Differentiating,

$$\begin{aligned} \frac{\partial H_N(q; \rho)}{\partial q_{x'}} &= - \sum_{i \in \mathcal{I}: x' \in X_i} c_i |X_i|^{1-\rho} \frac{1}{\rho-1} \bar{q}_i^{\rho-1} (e_{x'}^T q)^{1-\rho} \\ &\quad + \sum_{i \in \mathcal{I}: x' \in X_i} c_i |X_i|^{1-\rho} \frac{1}{\rho-2} \bar{q}_i^{\rho-2} \sum_{x'' \in X_i} (e_{x''}^T q)^{2-\rho} \\ &\quad - \sum_{i \in \mathcal{I}: x' \in X_i} c_i \frac{1}{|X_i|} \frac{1}{(\rho-2)(\rho-1)}. \end{aligned}$$

Differentiating again,

$$\begin{aligned} \frac{\partial^2 H_N(q; \rho)}{\partial q_{x'} \partial q_{x''}} &= \delta_{x', x''} \sum_{i \in \mathcal{I}: x' \in X_i} c_i |X_i|^{1-\rho} \bar{q}_i^{\rho-1} q_{x'}^{-\rho} \\ &\quad - \sum_{i \in \mathcal{I}: x', x'' \in X_i} c_i |X_i|^{1-\rho} \bar{q}_i^{\rho-2} q_{x'}^{1-\rho} \\ &\quad - \sum_{i \in \mathcal{I}: x', x'' \in X_i} c_i |X_i|^{1-\rho} \bar{q}_i^{\rho-2} q_{x''}^{1-\rho} \\ &\quad + \sum_{i \in \mathcal{I}: x', x'' \in X_i} c_i |X_i|^{1-\rho} \bar{q}_i^{\rho-3} \sum_{x''' \in X_i} q_{x'''}^{2-\rho}. \end{aligned}$$

Thus,

$$q_{x'} \left(\frac{\partial^2 H_N(q; \rho)}{\partial q_{x'} \partial q_{x''}} \right) q_{x''} = \sum_{i \in \mathcal{I}: x', x'' \in X_i} c_i |X_i|^{1-\rho} \bar{q}_i \left\{ \delta_{x', x''} \left(\frac{q_{x'}}{\bar{q}_i} \right)^{2-\rho} - \left(\frac{q_{x'}}{\bar{q}_i} \right)^{2-\rho} \left(\frac{q_{x''}}{\bar{q}_i} \right) - \left(\frac{q_{x''}}{\bar{q}_i} \right)^{2-\rho} \left(\frac{q_{x'}}{\bar{q}_i} \right) + \left(\frac{q_{x''}}{\bar{q}_i} \right) \left(\frac{q_{x'}}{\bar{q}_i} \right) \left(\sum_{x''' \in X_i} \left(\frac{q_{x'''}}{\bar{q}_i} \right)^{2-\rho} \right) \right\}.$$

Note that this equation also holds in the $\rho = 2$ and $\rho = 1$ cases. We can write this as

$$q_{x'} \left(\frac{\partial^2 H_N(q; \rho)}{\partial q_{x'} \partial q_{x''}} \right) q_{x''} = \sum_{i \in \mathcal{I}} c_i |X_i|^{1-\rho} \bar{q}_i e_x^T E_i^T m(q_i) E_i e_{x''},$$

where

$$\begin{aligned} m(q_i) &= \text{Diag}(q_i)^{2-\rho} - \text{Diag}(q_i)^{2-\rho} \iota q_i^T - q_i \iota^T \text{Diag}(q_i)^{2-\rho} + q_i \iota^T \text{Diag}(q_i)^{2-\rho} \iota q_i^T \\ &= (I - \iota q_i^T)^T \text{Diag}(q_i)^{2-\rho} (I - \iota q_i^T). \end{aligned}$$

The result immediately follows in the $\rho = 2$ case. For any $\rho \neq 2$,

$$\begin{aligned} m(q_i)^{\frac{1}{2-\rho}} &= (I - q_i \iota^T) \text{Diag}(q_i) (I - \iota q_i^T) \\ &= \text{Diag}(q) - q_i q_i^T - q_i q_i^T + q_i q_i^T \\ &= g^+(q_i). \end{aligned}$$

If $\rho < 2$, $H_N(q; \rho)$ is a bounded convex function on the relative interior of the simplex, and hence by theorem 10.3 of Rockafellar (1970) there is a unique extension to the simplex.

B.4 Proof of Lemma 3

First, note that if $\rho \geq 2$ and q_s does not have full support, then p_x will not have full support for the state x such that $e_x^T q_s = 0$, and we will have $D_\rho(p_x || p E_i^T q_i) = \infty$ for any i with $x \in X_i$, as required. For $\rho < 2$, continuity holds, and therefore both boundary cases are satisfied, provided the result holds for interior q_s .

To prove this claim, it is sufficient to show that, if all q_s are interior,

$$\sum_{i \in \mathcal{I}} c_i |X_i|^{1-\rho} \bar{q}_i^{\rho-1} \sum_{x \in X_i} (e_x^T q)^{2-\rho} D_\rho(p e_x || p E_i^T q_i) = -H_N(q) + \sum_{s \in S} (e_s^T p q) H_N(e_s^T p \text{Diag}(q)).$$

Using Lemma 2,

$$\sum_{s \in \mathcal{S}} \pi_s H_N(q_s) = \sum_{s \in \mathcal{S}: \pi_s > 0} \pi_s \sum_{i \in \mathcal{I}} c_i \bar{q}_{i,s} \frac{1}{|X_i|} \frac{1}{(\rho - 2)(\rho - 1)} \sum_{x \in X_i} \left\{ \left(\frac{e_x^T q_s}{\frac{1}{|X_i|} \bar{q}_{i,s}} \right)^{2-\rho} - 1 \right\}.$$

Using Bayes' rule, $\pi_s \bar{q}_{i,s} = \bar{q}_i \bar{p}_{i,s}$, where $\bar{p}_{i,s} = p E_i^T q_i$, and therefore

$$\begin{aligned} \sum_{s \in \mathcal{S}} \pi_s H_N(q_s) &= \sum_{i \in \mathcal{I}} c_i |X_i|^{1-\rho} \bar{q}_i^{\rho-1} \frac{1}{(\rho - 2)(\rho - 1)} \sum_{x \in X_i} (e_x^T q)^{2-\rho} \sum_{s \in \mathcal{S}: \pi_s > 0} \bar{p}_{i,s}^{\rho-1} (e_s^T p e_x)^{2-\rho} \\ &\quad - \sum_{i \in \mathcal{I}} c_i \bar{q}_i \frac{1}{(\rho - 2)(\rho - 1)}. \end{aligned}$$

Therefore,

$$-H_N(q) + \sum_{s \in \mathcal{S}} \pi_s H_N(q_s) = \sum_{i \in \mathcal{I}} c_i |X_i|^{1-\rho} \bar{q}_i^{\rho-1} \sum_{x \in X_i} (e_x^T q)^{2-\rho} D_\rho(p_x || p E_i^T q_i),$$

as required. The proof is essentially identical in the $\rho = 1$ and $\rho = 2$ cases.